



A fully polynomial-time approximation scheme for approximating a sum of random variables



Jian Li*, Tianlin Shi

Institute for Interdisciplinary Information Sciences, Tsinghua University, China

ARTICLE INFO

Article history:

Received 29 October 2013

Received in revised form

11 February 2014

Accepted 11 February 2014

Available online 18 February 2014

Keywords:

Threshold probability

Tail probability

Approximate counting

Counting knapsack

FPTAS

ABSTRACT

Given n independent integer-valued random variables X_1, X_2, \dots, X_n and an integer C , we study the fundamental problem of computing the probability that the sum $\mathbf{X} = X_1 + X_2 + \dots + X_n$ is at most C . We assume that each random variable X_i is implicitly given by an oracle \mathcal{O}_i , which given two input integers n_1, n_2 returns the probability of $n_1 \leq X_i \leq n_2$. We give the first deterministic fully polynomial-time approximation scheme (FPTAS) to estimate the probability up to a relative error of $1 \pm \epsilon$. Our algorithm is based on the technique for approximately counting knapsack solutions, developed in Gopalan et al. (2011).

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

We study the following fundamental problem. The input consists of n independent (not necessarily identically distributed) random integral variables X_1, \dots, X_n and an integer C . Our task is to compute the following probability value

$$F(C) = \Pr \left[\sum_{i=1}^n X_i \leq C \right]. \quad (1)$$

It is well known that computing $F(C)$ is #P-hard (see e.g., [9]). The hardness of computing $F(C)$ has an essential impact in the area of stochastic optimization as many problems generalize and/or utilize this basic problem in one way or another, thus inheriting the #P-hardness. Although we can sometimes use for example the linearity of expectation to bypass the difficulty of computing $F(C)$, more than often no such simple trick is applicable, especially in the context of risk-aware stochastic optimization where people usually pay more attention to the tail probability than the expectation.

Despite the importance of the problem, surprisingly, no approximation algorithm with provable multiplicative factor is known. We note that we can easily obtain an additive PRAS (polynomial-time randomized approximation scheme) for this problem via the

Monte-Carlo method: for each $i \in \{1, 2, \dots, n\}$, generate K independent samples $X_i^{(k)}$, $k = 1, 2, \dots, K$, according to the distribution of X_i , and then use the empirical average

$$\tilde{F}(C) = \frac{1}{K} \sum_{k=1}^K \mathbb{I} \left(\sum_{i=1}^n X_i^{(k)} \leq C \right)$$

as the estimation of $F(C)$, where $\mathbb{I}(\cdot)$ is the indicator function. It is easy to see that $\tilde{F}(C)$ is an unbiased estimator of $F(C)$. By standard Chernoff bound, one can see that with $K = \text{poly}(1/\epsilon)$ samples, the estimate is within an additive error ϵ from the true value with constant probability (see e.g., [14]). To get a reasonable multiplicative approximation factor (say a constant close to 1), we need to set the additive error at the order of $F(C)$. So the number of samples needs to be $\text{poly}(1/F(C))$, which can be exponentially large, when $F(C)$ is exponentially small. In certain application domains such as risk analysis, small probabilities (often associated with catastrophic losses) can be significant, thus demanding accurate estimations. Therefore, it is important to develop polynomial time approximation schemes for estimating such small probabilities.

Assumptions. Before presenting our main result, we need some notations and assumptions of the computation model. We assume that all random variables are discrete and the support of X_i , denoted as supp_i , is finite and consists of only integers. Without loss of generality, we can assume all X_i s are nonnegative (i.e., $\text{supp}_i \subseteq \mathbb{N}$) and $0 \in \text{supp}_i$ for all i . To see why this is without loss of generality, simply consider the equivalent problem of computing $\Pr[\sum_{i=1}^n (X_i - \min X_i) \leq C - \sum_{i=1}^n \min X_i]$, where $\min X_i$ is the

* Corresponding author.

E-mail addresses: lapordge@gmail.com, lijian83@mail.tsinghua.edu.cn (J. Li), stl501@gmail.com (T. Shi).

minimum value in supp_i . Under such an assumption, the problem is non-trivial only for $C > 0$. Moreover, we can assume that $\text{supp}_i \in [0, C + 1]$ for all i since we can place all mass in $[C + 1, \infty)$ at the single point $C + 1$, which does not affect the answer. The distribution of each random variable X_i is implicitly given by an oracle \mathcal{O}_i , which given two input value (n_1, n_2) returns the value $\Pr[n_1 \leq X_i \leq n_2]$ in constant time.

Our main result is a *fully polynomial-time approximation scheme* (FPTAS) for computing $F(C)$. For ease of notation, we use $(1 \pm \epsilon)F(C)$ to denote the interval $[(1 - \epsilon)F(C), (1 + \epsilon)F(C)]$. Let $\Delta = \prod_i \Pr[X_i = 0]$. Clearly Δ is a lower bound on the solution. Recall that we say there is an FPTAS for the problem, if for any positive constant $\epsilon \gtrsim 0$, there is an algorithm which can produce an estimate \tilde{F} with $\tilde{F} \in (1 \pm \epsilon)F(C)$ in $\text{poly}(n, \epsilon^{-1}, \log C, \log \frac{1}{\Delta})$ time (see e.g., [14]).

Theorem 1.1. *We are given n independent nonnegative integer-valued random variables X_1, \dots, X_n , a positive integer C , and a constant $\epsilon > 0$. Suppose that for all $i \in \{1, 2, \dots, n\}$, $\text{supp}_i \subseteq [0, C + 1]$, $0 \in \text{supp}_i$ and there is an oracle \mathcal{O}_i , which, upon two input integers (n_1, n_2) , returns the value $\Pr[n_1 \leq X_i \leq n_2]$ in constant time. There is an FPTAS for estimating $\Pr[\sum_{i=1}^n X_i \leq C]$ and the running time is $O(\frac{n^3}{\epsilon^2} \log(\frac{1}{\Delta})^2 \log C)$.*

Remark 1. For simplicity of presentation, we assume in the above theorem a computation model in which any real arithmetic can be performed with perfect accuracy in constant time and the probability values returned by the oracle are reals, also with perfect accuracy. In Section 2.3, we show how to implement our algorithm in a computation model where only bit operations are allowed and the oracles also return numerical values with finite precision. We show that the bit complexity of the algorithm is still a (somewhat larger) polynomial.

Remark 2. Note that, the oracle assumption is weaker than assuming the explicit representations of the distributions (i.e., listing the probability mass at every point). In fact, if the input is the explicit representations of the distributions, we can preprocess the input in linear time so that each oracle call to \mathcal{O}_i can be simulated in $O(\log |\text{supp}_i|)$ time. This can be done by computing the prefix sums $\Pr[X_i \leq x]$ for all $x \in \text{supp}_i$ in $O(|\text{supp}_i|)$ time. Then for each oracle call (n_1, n_2) , we use binary search to find out the smallest value $x_1 \in \text{supp}_i$ that is no smaller than n_1 and the largest value $x_2 \in \text{supp}_i$ that is no larger than n_2 in $O(\log |\text{supp}_i|)$ time. Therefore, $\Pr[n_1 \leq X_i \leq n_2]$ is the same as $\Pr[x_1 \leq X_i \leq x_2]$, which can be computed from the prefix sums in constant time.

1.1. Related work

There is a large body of work on estimating or upper/lower-bounding the distribution of the sum of independent random variables. See e.g., [1, 18, 12, 3, 13]. Those works are based on analytic numerical methods (e.g., Edgeworth expansion, saddle point method) which either require specific families of distributions and/or do not provide any provable multiplicative approximation guarantees.

Our problem is a generalization of the counting knapsack problem. For the counting knapsack problem, Morris and Sinclair [15] obtained the first FPRAS (fully polynomial-time randomized approximation scheme) based on the Markov Chain Monte-Carlo (MCMC) method. Dyer [4] provided a completely different FPRAS based on dynamic programming. The first deterministic FPTAS is obtained by Gopalan et al. [7] (see also the journal version [19]).

Our problem is also closely related to the threshold probability maximization problem (see a general formulation in [11]). In

this problem, we are given a ground set of items. Each feasible solution to the problem is a subset of the elements satisfying some property (this includes problems such as shortest path, minimum spanning tree, and minimum weight matching). Each element b is associated with a random weight X_b . Our goal is to find a feasible set S such that $\Pr[\sum_{b \in S} X_b \leq C]$ is maximized, for a given threshold C . There is a large body of literature on the threshold probability maximization problem, especially for specific combinatorial problems and/or special distributions. For example, Nikolova, Kelner, Brand and Mitzenmacher [17] studied the corresponding shortest path version for Gaussian, Poisson and exponential distributions. Nikolova [16] extended this result to an FPTAS for any problem with Gaussian distributions, if the deterministic version of the problem has a polynomial-time algorithm. The minimum spanning tree version with Gaussian distributed edges has also been studied in [5]. For general discrete distributions, Li and Deshpande [10] obtained an additive PTAS if the deterministic version of the problem can be solved exactly in pseudo-polynomial time. Very recently, Li and Yuan [11] further generalized this result to the class of problems for which the multi-objective deterministic version admits a PTAS.

Our problem is also closely related to the fixed set version of the stochastic knapsack problem. In this problem, we are given a knapsack of capacity C and a set of items with random sizes and profits. Their goal is to find a set of items with maximum total profit subject to the constraint that the overflow probability is at most a given parameter γ . Kleinberg, Rabani and Tardos [9] first considered the problem with Bernoulli-type distributions and provided a polynomial-time $O(\log 1/\gamma)$ -approximation. Better results are known for specific distributions, such as exponentially distributions [6], Gaussian distributions [8, 16]. For general discrete distributions, bi-criteria additive PTASes (i.e., the overflow probability constraint may be violated by an additive factor ϵ for any constant $\epsilon > 0$) are known via different techniques [2, 10, 11].

2. Algorithm

Our algorithm is based on dynamic programming. In Section 2.1, we provide the recursion of the dynamic program, which is largely based on the idea developed in [7, 19], with some necessary adaptations. However, since the support of each random variable can be exponentially large, it is not immediately clear how the recursion can be implemented efficiently given the oracles. We address this issue in Section 2.2. In Section 2.3, we analyze the bit complexity of our algorithm.

2.1. The dynamic program

We first notice that $\Pr[\sum_{j=1}^i X_j \leq C]$, for any $i \in \{1, 2, \dots, n\}$, is a nondecreasing function of C . We consider its inverse function $\tau(i, a) : \{1, 2, \dots, n\} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{N} \cup \{\pm\infty\}$, which is defined to

$$\tau(i, a) = \begin{cases} \min \left\{ C \mid C \geq 0 \text{ and } \Pr \left[\sum_{j=1}^i X_j \leq C \right] \geq a \right\}, & 0 < a \leq 1; \\ +\infty, & a > 1; \\ -\infty, & a = 0. \end{cases}$$

It is easy to see that $\tau(i, a)$ is nondecreasing in a . The following simple lemma is needed. We omit the proof, which is straightforward.

Lemma 2.1. *Both of the following statements hold true.*

1. $\Pr \left[\sum_{i=1}^n X_i \leq C \right] = \max \{ a : \tau(n, a) \leq C \}$.
2. $\tau(i, a) = 0$ for any $i \in \{1, 2, \dots, n\}$ and $a \leq \Delta$, where $\Delta = \prod_i \Pr[X_i = 0]$.

The following recursion is very important to us. For $x \in \text{supp}_i$, we use $p_i(x)$ as a shorthand notation for $\Pr[X_i = x]$.

Lemma 2.2. Let $\gamma_i : \text{supp}_i \rightarrow \mathbb{R}_{\geq 0}$ denote a function. The following recurrence holds:

$$\tau(i, a) = \min_{\gamma_i} \max_{x \in \text{supp}_i} \{ \tau(i-1, \gamma_i(x)) + x \},$$

for any $i \in \{1, 2, \dots, n\}$ and $a \geq 0$, (2)

where the minimum is taken among all functions γ_i such that

$$\sum_{x \in \text{supp}_i} \gamma_i(x) p_i(x) \geq a. \tag{3}$$

Intuitively, $\gamma_i(x)$ represents the value $\Pr[\sum_{j=1}^{i-1} X_j \leq \tau(i, a) - x]$. By enforcing (3), we make sure that $\Pr[\sum_{j=1}^i X_j \leq \tau(i, a)] \geq a$ so that the definition of $\tau(i, a)$ is met. The formal proof is as follows.

Proof. We first fix some $0 < a \leq 1$ and prove that for any $\gamma_i : \text{supp}_i \rightarrow [0, 1]$, the quantity $C' := \max_{x \in \text{supp}_i} \{ \tau(i-1, \gamma_i(x)) + x \} \geq \tau(i, a)$. Due to the independence, we can see that

$$\Pr \left[\sum_{j=1}^i X_j \leq C' \right] = \sum_{x \in \text{supp}_i} p_i(x) \cdot \Pr \left[\sum_{j=1}^{i-1} X_j \leq C' - x \right]. \tag{4}$$

By the definition of C' , we know that $C' - x \geq \tau(i-1, \gamma_i(x))$ for all $x \in \text{supp}_i$, and hence $\Pr[\sum_{j=1}^{i-1} X_j \leq C' - x] \geq \gamma_i(x)$. This gives $\Pr[\sum_{j=1}^i X_j \leq C'] \geq \sum_{x \in \text{supp}_i} \gamma_i(x) p_i(x) \geq a$. Therefore, we have $C' \geq \tau(i, a)$. If $a > 1$, by (3) we know that $\gamma_i(x') > 1$ for some $x' \in \text{supp}_i$ and therefore $C' \geq \tau(i-1, \gamma_i(x')) = \infty$. If $a = 0$, $\tau(i, a)$ is defined to be $-\infty$ and we trivially have $C' \geq \tau(i, a)$. In all above cases, we have $C' \geq \tau(i, a)$.

Now, we prove the reverse direction. In particular, we prove that there exists a choice of $\gamma_i := \gamma_i^*$ such that $C^* := \max_{x \in \text{supp}_i} \{ \tau(i-1, \gamma_i^*(x)) + x \} \leq \tau(i, a)$. First assume that $0 < a \leq 1$. Let $\gamma_i^*(x) = \Pr[\sum_{j=1}^{i-1} X_j \leq \tau(i, a) - x]$. It is easy to see that γ_i^* satisfies (3), if we let $C' = \tau(i, a)$ in (4) and use the definition of $\tau(i, a)$ on the LHS. Therefore, $\tau(i-1, \gamma_i^*(x)) \leq \tau(i, a) - x$ for all $x \in \text{supp}_i$, which leads to $\max_{x \in \text{supp}_i} \{ \tau(i-1, \gamma_i^*(x)) + x \} \leq \tau(i, a)$. If $a = 0$, let $\gamma_i^*(x) = 0$ for all $x \in \text{supp}_i$. It is easy to see that $C^* \leq \tau(i, a)$ since $C^* = -\infty$ by the definition of C^* . If $a > 1$, $\tau(i, a) = +\infty$ and $C^* \leq \tau(i, a)$ is also trivially true. This completes the proof. \square

For ease of notation, we define function $\lambda_i : \text{supp}_i \rightarrow \mathbb{R}_{\geq 0}$ as $\lambda_i(x) = \gamma_i(x) p_i(x) / a$. Hence, the recursion in Lemma 2.2 can be rewritten as:

$$\tau(i, a) = \min_{\lambda_i} \max_{x \in \text{supp}_i} \left\{ \tau \left(i-1, \frac{\lambda_i(x)}{p_i(x)} a \right) + x \right\}$$

subject to $\sum_{x \in \text{supp}_i} \lambda_i(x) \geq 1$. (5)

Since the second argument a is a continuous variable, it is not clear how recursion (5) can be efficiently evaluated. To overcome this issue, we discretize the above recursion as follows: let

$$Q = 1 + \frac{\epsilon}{n}.$$

Note that Q is slightly larger than 1. Recall that $\Delta = \prod_i \Pr[X_i = 0]$, which is a trivial lower bound on $F(C)$. Let $s = \lceil \log_Q \frac{1}{\Delta} \rceil = O(\frac{n}{\epsilon} \log \frac{1}{\Delta})$.

We define recursively the function $T : \{1, 2, \dots, n\} \times \{0, 1, 2, \dots, s\} \rightarrow \mathbb{Z} \cup \{\pm\infty\}$ as follows: the base case is defined as

$$T(1, j) = \tau(1, Q^{-j}), \quad \text{for all } j \in [s]. \tag{6}$$

Algorithm 1: The Dynamic Program for Computing (1)

```

1 Let  $Q = 1 + \epsilon/n$  and  $s = \lceil \log_Q \frac{1}{\Delta} \rceil$ ;
2 Query oracle  $\mathcal{O}_1$  to obtain  $T(1, j) = \tau(1, Q^{-j})$  for all  $j \in [s]$ ;
3 for  $i = 2 \rightarrow n$  do
4   for  $j = 0 \rightarrow s$  do
5      $T(i, j) \leftarrow$ 
        $\left[ \begin{array}{l} \min_{\lambda_i} \max_{x \in \text{supp}_i} T(i-1, j + \lfloor \log_Q \frac{p_i(x)}{\lambda_i(x)} \rfloor) + x, \\ \text{subject to } \sum_{x \in \text{supp}_i} \lambda_i(x) \geq 1. \end{array} \right.$ 
6  $j^* \leftarrow \min\{j : T(n, j) \leq C\}$ . Return  $\tilde{F}(C) := Q^{-j^*+1}$ ;
```

Note that each $\tau(1, Q^{-j})$ can be computed with $O(\log C)$ calls to \mathcal{O}_1 . For $i \geq 2$ and $j \in [s]$, define

$$T(i, j) = \min_{\lambda_i} \max_{x \in \text{supp}_i} \left\{ T \left(i-1, j + \left\lfloor \log_Q \frac{p_i(x)}{\lambda_i(x)} \right\rfloor \right) + x \right\}$$

subject to $\sum_{x \in \text{supp}_i} \lambda_i(x) \geq 1$. (7)

To make analysis easier (mainly to avoid tedious case-by-case study near the boundary), we extend the domain of $T(i, j)$ to $\{1, 2, \dots, n\} \times \mathbb{Z} \cup \{+\infty\}$ such that for all $i \in \{1, 2, \dots, n\}$, we define $T(i, j) = 0$ for $j > s$ and $T(i, j) = +\infty$ for $j < 0$. Furthermore, we define $T(i, +\infty) = -\infty$. So if $\lambda_i(x) = 0$, $T(i-1, j + \lfloor \log_Q \frac{p_i(x)}{\lambda_i(x)} \rfloor) = T(i-1, +\infty) = -\infty$. Comparing (7) and (5), the similarity suggests that $T(i, j)$ is an approximate version of $\tau(i, Q^{-j})$. The next lemma formalizes this idea.

Lemma 2.3. For all $i \in \{1, 2, \dots, n\}$ and $j \in \mathbb{Z} \cup \{\infty\}$, we have that $\tau(i, Q^{-j}) \leq T(i, j) \leq \tau(i, Q^{-(j-i)})$.

Proof. We prove this by induction. The base case is trivially true by the definition of $T(1, j)$, even for $j < 0$ and $j > s$. Now assume that the statement is true for $i-1$ ($i \geq 2$) and all $j \in \mathbb{Z} \cup \{+\infty\}$. We prove it is also true for i and all $j \in \mathbb{Z} \cup \{+\infty\}$. By the induction hypothesis, we have that

$$T(i-1, \lfloor j + \log_Q \frac{p_i(x)}{\lambda_i(x)} \rfloor) \leq \tau(i-1, Q^{-(\lfloor j + \log_Q \frac{p_i(x)}{\lambda_i(x)} \rfloor - i + 1)})$$

$$\leq \tau(i-1, \frac{\lambda_i(x)}{p_i(x)} Q^{-(j-i)}) \quad \text{and}$$

$$T(i-1, \lfloor j + \log_Q \frac{p_i(x)}{\lambda_i(x)} \rfloor) \geq \tau(i-1, Q^{-\lfloor j + \log_Q \frac{p_i(x)}{\lambda_i(x)} \rfloor})$$

$$\geq \tau(i-1, \frac{\lambda_i(x)}{p_i(x)} Q^{-j}).$$

Note that the above inequalities hold even $j + \log_Q \frac{p_i(x)}{\lambda_i(x)}$ is negative or larger than s , or $\lambda_i(x) = 0$. Taking the maximum over x and the minimum over λ_i does not change the direction of the inequalities. Combining the above inequalities with (5), we complete the proof. \square

With this lemma, we can approximate the recursion (5) by solving the recursion (7). The pseudocode of the dynamic program is provided in Algorithm 1. Note that it is still not clear how to implement (7) in polynomial time, as the number of possible functions λ_i can be infinite. We address this issue in the next section. Assuming (7) can be implemented efficiently, we can show the output of the algorithm is a good approximation of the true probability with the following lemma.

Algorithm 2: Efficient Implementation of Recursion (7)

```

1 Initially, let  $\mathcal{L} = 0, \mathcal{R} = n(C + 1)$ ;
2 while  $\mathcal{R} > \mathcal{L}$  do
3    $\mathcal{T} \leftarrow \lfloor (\mathcal{L} + \mathcal{R})/2 \rfloor$ ;
4    $P_i(0) \leftarrow \Pr[X_i \in (\mathcal{T} - \infty, \mathcal{T} - T(i - 1, 0))]$  (via oracle  $\mathcal{O}_i$ );
5   for  $m = 1, 2, \dots, s + 1$  do
6      $P_i(m) \leftarrow \Pr[X_i \in (\mathcal{T} - T(i - 1, m - 1), \mathcal{T} - T(i - 1, m))]$ 
       (via oracle  $\mathcal{O}_i$ );
7    $Z = \sum_{m=0}^{s+1} Q^{j+s+1-m} P_i(m)$ ;
8   if  $Z \geq Q^{s+1}$  then
9      $\mathcal{R} \leftarrow \mathcal{T}$ ;
10  else
11     $\mathcal{L} \leftarrow \mathcal{T} + 1$ ;
12 Return  $T(i, j) := \mathcal{L}$ ;
```

Lemma 2.4. The output $\tilde{F}(C)$ of Algorithm 1 is a $(1 \pm \epsilon)$ -approximation of $F(C)$.

Proof. From the choice of j^* , we know that $T(n, j^*) \leq C < T(n, j^* - 1)$. According to Lemma 2.3, we have $\tau(n, Q^{-j^*}) \leq C < \tau(n, Q^{-j^*+1})$. Therefore, $F(C) = \Pr[\sum_{i=1}^n X_i \leq C] \in [Q^{-j^*}, Q^{-j^*+1}]$. By outputting $\tilde{F}(C) = Q^{-j^*+1}$, the approximation ratio can be bounded as $1 - \epsilon \leq Q^{-n} \leq \tilde{F}(C)/F(C) \leq Q \leq 1 + \epsilon$. \square

2.2. An efficient implementation using binary search

In this subsection, we show how to implement the recursion (7) in polynomial time. Suppose that we have already computed $T(i', j')$ for all $i' < i$ and $0 \leq j' \leq s$ and we are trying to compute the value $T(i, j)$. Our approach is based on a binary search over the range of $T(i, j)$. We maintain an interval $[\mathcal{L}, \mathcal{R}]$ such that $T(i, j) \in [\mathcal{L}, \mathcal{R}]$ throughout the algorithm. In each iteration, we make a guess $\mathcal{T} = \lfloor (\mathcal{L} + \mathcal{R})/2 \rfloor$ and decide whether $T(i, j) \leq \mathcal{T}$ using the criterion in Lemma 2.5. It is not immediately clear how we can check efficiently whether the criterion is met. We address this issue in Lemma 2.7. The pseudocode is provided in Algorithm 2.

Lemma 2.5. Suppose that we are computing $T(i, j)$ and \mathcal{T} is the current guess. We define that for all $x \in \text{supp}_i$,

$$\lambda'_i(x) = \max \left\{ z \mid z \in \mathbb{R}_{\geq 0} \text{ and } T(i - 1, j + \lfloor \log_Q \frac{p_i(x)}{z} \rfloor) + x \leq \mathcal{T} \right\}.$$

Then, $T(i, j) \leq \mathcal{T}$ if and only if $\sum_{x \in \text{supp}_i} \lambda'_i(x) \geq 1$.

Proof. Suppose $\sum_{x \in \text{supp}_i} \lambda'_i(x) \geq 1$. This means that there exists a choice of $\lambda_i := \lambda'_i$ such that

$$\max_{x \in \text{supp}_i} \left\{ T(i - 1, j + \lfloor \log_Q \frac{p_i(x)}{\lambda_i(x)} \rfloor) + x \right\} \leq \mathcal{T},$$

and therefore $T(i, j) \leq \mathcal{T}$ by (7).

On the other hand, assume $T(i, j) \leq \mathcal{T}$. Suppose that $T(i, j)$ is achieved by $\lambda_i := \lambda_i^*$ in (7). We know that $T(i - 1, j + \lfloor \log_Q \frac{p_i(x)}{\lambda_i^*(x)} \rfloor) + x \leq T(i, j) \leq \mathcal{T}$ for any $x \in \text{supp}_i$. By definition of λ'_i , we see that $\lambda_i^*(x) \leq \lambda'_i(x)$, for all $x \in \text{supp}_i$. Therefore, $\sum_{x \in \text{supp}_i} \lambda'_i(x) \geq \sum_{x \in \text{supp}_i} \lambda_i^*(x) \geq 1$. \square

From Lemma 2.5, we can see that if $|\text{supp}_i|$ is bounded by a polynomial, we can decide whether $T(i, j) \leq \mathcal{T}$ in polynomial time. However, $|\text{supp}_i|$ can be exponentially large. We resolve this issue by showing that $\sum_{x \in \text{supp}_i} \lambda'_i(x)$ can be computed using at most $O(s)$ calls to \mathcal{O}_i . For this purpose, we divide the support of X_i into $s + 3$ segments such that the sum of $\lambda'_i(x)$ values over each segment can be computed efficiently. In particular, we divide \mathbb{N} into segments

$$\text{Seg}_0 = (-\infty, \mathcal{T} - T(i - 1, 0)) \cap \mathbb{N},$$

$$\text{Seg}_k = (\mathcal{T} - T(i - 1, k - 1), \mathcal{T} - T(i - 1, k)) \cap \mathbb{N},$$

for all $k = 1, \dots, s + 1$,

$$\text{Seg}_{s+2} = (\mathcal{T} - T(i - 1, s + 1), \infty) \cap \mathbb{N}.$$

Let us first investigate the last segment.

Lemma 2.6. Suppose $x \in \text{Seg}_{s+2}$ and $\lambda'_i(x)$ is defined as in Lemma 2.5. Then, we have that $\lambda'_i(x) = 0$.

Proof. First, recall that $T(i - 1, s + 1) = 0$. For any $x > \mathcal{T}$, the only possible nonnegative value z that can make $T(i - 1, j + \lfloor \log_Q (\frac{p_i(x)}{z}) \rfloor) + x \leq \mathcal{T}$ is 0, since $T(i - 1, +\infty) = -\infty$. \square

For each other segment, $\lambda'_i(x)$ can be computed using the following lemma.

Lemma 2.7. Suppose $x \in \text{Seg}_m$ for some $m \in [0, s + 1]$ and $\lambda'_i(x)$ is defined as in Lemma 2.5. Then, we have that

$$\lambda'_i(x) = Q^{j-m} p_i(x). \quad (8)$$

Proof. Fix any $x \in \text{Seg}_m$ for some $m \in [0, s + 1]$. Let $z > 0$ be such that

$$j + \log_Q \left(\frac{p_i(x)}{z} \right) = m.$$

Now, we show $\lambda'_i(x) = z$. Since $x \in \text{Seg}_m$, we can see that $\mathcal{T} - x \in [T(i - 1, m), T(i - 1, m - 1))$. Therefore, we have that

$$T(i - 1, j + \lfloor \log_Q \left(\frac{p_i(x)}{z} \right) \rfloor) = T(i - 1, m) \leq \mathcal{T} - x.$$

Now, we show any value z' larger than z does not satisfy the above inequality. This is because $T(i - 1, j + \lfloor \log_Q \left(\frac{p_i(x)}{z'} \right) \rfloor)$ is at least $T(i - 1, m - 1)$ which is greater than $\mathcal{T} - x$. Since z is the largest value that makes the inequality true, we have $\lambda'_i(x) = z$. \square

As a result, it is sufficient to ask the \mathcal{O}_i for the probability values $P_i(m) = \Pr[X_i \in \text{Seg}_m]$ and we can compute $\sum_{x \in \text{Seg}_m} \lambda'_i(x)$ simply by

$$\sum_{x \in \text{Seg}_m} \lambda'_i(x) = Q^{j-m} P_i(m). \quad (9)$$

Together with Lemma 2.5, (9) shows that the criterion $T(i, j) \leq \mathcal{T}$ is equivalent to $\sum_{m=0}^{s+1} Q^{j-m} P_i(m) \geq 1$. Multiplying Q^{s+1} on both sides, we derive Line 8 of Algorithm 2, which is used to decide the criterion $T(i, j) \leq \mathcal{T}$.

Lemma 2.8. Suppose that the oracle can answer each query in constant time. The optimization over all λ_i and x in Eq. (7) can be done in time $O(\frac{n}{\epsilon} \log \frac{1}{\Delta} \log C)$, where $\Delta = \prod_i \Pr[X_i = 0]$.

Proof. The binary search takes $\log C$ iterations, while in each iteration we make at most $O(s)$ oracle calls. So the implementation of the recursion step runs in $O(s \log C) = O(\frac{n}{\epsilon} \log \frac{1}{\Delta} \log C)$ since $s = O(\frac{n}{\epsilon} \log \frac{1}{\Delta})$. \square

The number of states in the dynamic program is $O(ns)$. Combined with Lemma 2.8, the total running time of Algorithm 1, which calls Algorithm 2 as a subroutine, is $O(\frac{n^3}{\epsilon^2} \log(\frac{1}{\Delta})^2 \log C)$. This completes the proof of Theorem 1.1.

2.3. The bit complexity of the algorithm

Our results have been stated under the assumption that any real arithmetic operation has unit cost. Now we consider the computation model where only bit operations are allowed. Moreover, the oracles also return numerical values that are encoded in bits. To model this, we assume that each query to the oracle contains an additional integer ℓ (called precision parameter), which specifies the number of bits the oracle should return. Upon such a query, the oracle returns the first (i.e., the most significant) ℓ bits of the answer, which is always between 0 and 1. This can model the case where the oracle is given by a mathematical formula and there is a numerical algorithm that can compute the answer to arbitrary precision. Our goal is to show that a slightly modified algorithm still preserves a $(1 \pm \epsilon)$ -approximation, while the bit complexity is also $\text{poly}(n, \epsilon^{-1}, \log C, \log \frac{1}{\Delta})$.

First, we use the oracle to figure out a close estimate of $\log 1/\Delta$, since it determines s , which relates to the table size of the dynamic problem. For each θ_i , we repeatedly query the value $\Pr[X_i = 0]$ with precision $\ell = 1, 2, 3, \dots$ until we get the first nonzero bit. It is easy to see that we stop with $\ell_i = \lceil -\log \Pr[X_i = 0] \rceil + 1$ steps. It then suffices for our algorithm to use an upper bound of $\log 1/\Delta$, which we choose to be $\sum_i \ell_i$.

In the sequel, we run Algorithms 1 and 2 as its subroutine with a slightly different value of $Q = 1 + 2^{-\ell_q}$, where $\ell_q = \lceil \log_2 \frac{n}{\epsilon} \rceil + 2$, so that $Q \leq 1 + \frac{\epsilon}{4n}$. Notice that $\mathcal{L}, \mathcal{R}, \mathcal{T}$ can be encoded in $O(\log C)$ bits and with $O(\log \frac{n}{\epsilon})$ bits, Q has exact representation. When we multiply a number of a bits and a number of b bits, we use $a + b$ bits for the result so that no precision is lost. Hence, any $Q^{j+s+1-m}$ ($0 \leq m \leq s+1$) used in the algorithm can be encoded in $O(s \log \frac{n}{\epsilon})$ bits.

The problematic part is the answer $P_i(m)$ returned by the oracle (in Steps 4 and 7 of Algorithm 2) is of finite precision. Whenever Algorithm 2 attempts to make an oracle call, e.g., $\Pr[n_1 \leq X_i \leq n_2]$, we set the precision parameter to be L , which we will determine later, and call θ_i , which returns a truncated probability $\widehat{P}_i(m)$ with a truncation error $P_i(m) - \widehat{P}_i(m) \in [0, 2^{-L+1}]$. To distinguish from the original version of Algorithm 2, we call the new version with finite precision *the bit version*. We use $\widehat{T}(i, j)$ to denote the $T(i, j)$ value obtained from the bit version of Algorithm 2. To account for this truncation error, we need a modified version of Lemma 2.3:

Lemma 2.9. *Let the precision parameter be L for each oracle call. For all $i \in \{1, 2, \dots, n\}$ and $j \in \mathbb{Z} \cup \{+\infty\}$, we have*

$$\tau(i, Q^{-j}) \leq \widehat{T}(i, j) \leq \tau(i, (1 + \eta)^i Q^{-(j-i)}), \quad (10)$$

where $\eta = (Q^{s+1} - Q^{-1}) / (Q - 1) \cdot 2^{-(L-1)}$.

Proof. The statement is trivial for $j > s$ and $j < 0$. For $j \in [s]$, we prove the lemma by induction on i .

We first prove the induction step. The proof for the base case $i = 1$ is very similar and we present it at the end. For $i \geq 2$, suppose in the dynamic program, we have computed the $\widehat{T}(i-1, j)$, $\forall j \in \mathbb{Z}$. Assume that the statement of the lemma is true for $i-1$ ($i \geq 2$) and $j \in \mathbb{Z}$. Now, we prove it for i and $j \in \mathbb{Z}$. Fix a particular $0 \leq j \leq s$. Imagine that we run two copies of Algorithm 2 simultaneously: one copy \mathcal{A}_1 is the original version with infinite precision, and the other \mathcal{A} is the bit version. We note that both \mathcal{A}_1 and \mathcal{A} use $\widehat{T}(i-1, j)$ values. So the output of \mathcal{A}_1 may not be the exact value of $T(i, j)$. Clearly, the execution flows of both \mathcal{A}_1 and \mathcal{A} would be the same in the beginning (i.e., $\mathcal{L}, \mathcal{R}, \mathcal{T}$ values are the same). If the execution flows are the same throughout, the outputs of \mathcal{A}_1 and \mathcal{A} are the same. Otherwise, we consider the snapshot when the deviation first occurs and the condition in Line 8 returns different boolean values between \mathcal{A}_1 and \mathcal{A} . Recall that $Z = \sum_{m=0}^{s+1} Q^{j+s+1-m} P_i(m)$ and let $\widehat{Z} = \sum_{m=0}^{s+1} Q^{j+s+1-m} \widehat{P}_i(m)$. It is easy to see that $\widehat{Z} < Q^{s+1}$

$\leq Z$. So the only possibility is that \mathcal{A}_1 jumps to Line 9 and \mathcal{A} goes to Line 11. Let T_1 be output of \mathcal{A}_1 . The above argument shows that $\widehat{T}(i, j) \geq T_1$.

Now we show $T_1 \geq \tau(i, Q^{-j})$. By (7) and Lemma 2.5 (with $T(i-1, \cdot)$ replaced by $\widehat{T}(i-1, \cdot)$), T_1 is the solution of the following optimization problem:

$$T_1 = \min_{\lambda_i} \max_{x \in \text{supp}_i} \left\{ \widehat{T}(i-1, j + \lfloor \log_Q \left(\frac{P_i(x)}{\lambda_i(x)} \right) \rfloor) + x \right\},$$

$$\text{subject to } \sum_{x \in \text{supp}_i} \lambda_i(x) \geq 1.$$

Using the same proof as in Lemma 2.3, we can show that $T_1 \geq \tau(i, Q^{-j})$. Hence, we have one direction of the lemma: $\widehat{T}(i, j) \geq \tau(i, Q^{-j})$.

The proof for the other direction is similar. This time, we run \mathcal{A} simultaneously with a slightly modified version \mathcal{A}_2 of Algorithm 2. Suppose that the oracles in \mathcal{A}_2 can return values with infinite precision. The goal of \mathcal{A}_2 is to solve the following optimization problem:

$$T_2 = \min_{\lambda_i} \max_{x \in \text{supp}_i} \left\{ \widehat{T}(i-1, j + \lfloor \log_Q \left(\frac{P_i(x)}{\lambda_i(x)} \right) \rfloor) + x \right\},$$

$$\text{subject to } \sum_{x \in \text{supp}_i} \lambda_i(x) \geq 1 + \eta. \quad (11)$$

We immediately can see that T_2 is an approximate value of $\tilde{\tau}(i, Q^{-(j-i)})$, which is defined as

$$\tilde{\tau}(i, Q^{-(j-i)}) = \min_{\lambda_i} \max_{x \in \text{supp}_i} \left\{ \tau(i-1, \frac{\lambda_i(x)}{P_i(x)} Q^{-(j-i)}) + x \right\},$$

$$\text{subject to } \sum_{x \in \text{supp}_i} \lambda_i(x) \geq 1 + \eta. \quad (12)$$

Using the same argument as in Lemma 2.3, we can show that

$$T_2 \leq \tilde{\tau}(i, (1 + \eta)^{i-1} Q^{-(j-i)}). \quad (13)$$

Furthermore, comparing (12) with (5), we can see the connection between τ and $\tilde{\tau}$ as follows:

$$\tilde{\tau}(i, (1 + \eta)^{i-1} Q^{-(j-i)}) = \tau(i, (1 + \eta)^i Q^{-(j-i)}). \quad (14)$$

Using the same argument as in Section 2.2, the optimization problem (11) can be implemented by a binary search procedure almost identical to Algorithm 2 except that \mathcal{A}_2 compares Z with $(1 + \eta)Q^{s+1}$ in Line 8. Notice that η has been chosen carefully such that the inequality $Z \geq (1 + \eta)Q^{s+1}$ implies $\widehat{Z} \geq Q^{s+1}$. Using a similar execution-flow analysis, we can see that the deviation occurs when $Q^{s+1} \leq \widehat{Z} < Z < (1 + \eta)Q^{s+1}$ and consequently we have that $T_2 \geq \widehat{T}(i, j)$. By (13) and (14), we have that $\widehat{T}_2 \leq \tau(i, (1 + \eta)^{i-1} Q^{-(j-i)})$.

Now we prove the base case $i = 1$. The proof is very similar to the induction step. Recall that in Algorithm 1, we use binary search to compute the value $T(1, j)$ exactly. The binary search maintains an interval $[\mathcal{L}, \mathcal{R}]$ such that $T(1, j) \in [\mathcal{L}, \mathcal{R}]$ throughout. In each iteration, we make a guess \mathcal{T} , compare it with Q^{-j} , and update $\mathcal{L} \leftarrow \mathcal{T} + 1$ if $\Pr[X_1 \leq \mathcal{T}] < Q^{-j}$, or $\mathcal{R} \leftarrow \mathcal{T}$ otherwise. In the bit version, Q^{-j} and $\Pr[X_1 \leq \mathcal{T}]$ may not be exact. Nevertheless, we can still get approximation $\widehat{T}(1, j)$, which is $T(1, j)$ of finite precision using a slightly modified binary search program \mathcal{B} . Denote the truncated probability from the oracle θ_1 as $\Pr[X_1 \leq \mathcal{T}]$. \mathcal{B} follows the previous binary search scheme but replaces the decision rule $\Pr[X_1 \leq \mathcal{T}] < Q^{-j}$ by $Q^j \cdot \Pr[X_1 \leq \mathcal{T}] < 1$. We first show the output of \mathcal{B} , $\widehat{T}(1, j)$ is at least $\tau(1, Q^{-j})$. Imagine another copy of this binary search program \mathcal{B}_1 running simultaneously but with infinite precision. Obviously, \mathcal{B}_1 computes the exact value $T(1, j) = \tau(1, Q^{-j})$. If the execution flows of \mathcal{B} and \mathcal{B}_1 are the

same throughout, then the output would be the same: $\widehat{T}(1, j) = T(1, j)$. Otherwise, by considering the snapshot when the deviation first occurs, i.e., $Q^j \cdot \Pr[X_1 \leq \mathcal{T}] < 1 \leq Q^j \cdot \widehat{\Pr}[X_1 \leq \mathcal{T}]$. It is not hard to see that $\widehat{T}(1, j) \geq T(1, j) = \tau(1, Q^{-j})$. Next, we show the upper bound by proving $\widehat{T}(1, j) \leq \tau(1, (1 + \eta)Q^{-j}) \leq \tau(1, (1 + \eta)Q^{1-j})$. Imagine a third copy of this binary search program \mathcal{B}_2 running simultaneously, but with the decision rule replaced by $Q^j \cdot \widehat{\Pr}[X_1 \leq \mathcal{T}] < 1 + \eta$. By the choice of η , the same execution-flow analysis gives that $\widehat{T}(1, j) \leq \tau(1, (1 + \eta)Q^{-j})$. This concludes the proof of the base case. \square

When all $\widehat{T}(n, j)$ values have been computed (for $j \in [0, s]$), the bit version of Algorithm 1 returns Q^{-j^*+1} up to precision ℓ_{ans} . The following lemma shows the returned $\widetilde{F}(C)$ is an $(1 \pm \epsilon)$ -approximation of $F(C)$.

Lemma 2.10. *By choosing $Q = 1 + 2^{-l_q} \leq 1 + \frac{\epsilon}{4n}$ for $\ell_q = \lceil \log_2 \frac{n}{\epsilon} \rceil + 2$, $L \geq \log_2 \frac{32n^2 Q}{\epsilon^2} + s + 1$ and $\ell_{\text{ans}} = s + \log_2 \frac{Q}{\epsilon}$, the above modified algorithm can produce a $(1 \pm \epsilon)$ -approximation of $F(C)$.*

Proof. According to Lemma 2.9, we have $\tau(n, Q^{-j^*}) \leq C \leq \tau(n, Q^{n-j^*+1}(1 + \eta)^n)$, and hence $F(C) \in [Q^{-j^*}, Q^{n-j^*+1}(1 + \eta)^n]$. On the other hand, the returned $\widetilde{F}(C)$ by our bit version of the algorithm satisfies $\widetilde{F}(C) \in (Q^{-j^*+1} - 2^{-\ell_{\text{ans}}+1}, Q^{-j^*+1}]$. By our choice of Q (or ℓ_q), L and ℓ_{ans} , it is not hard to verify that $\widetilde{F}(C)/F(C) \in [1 - \epsilon, 1 + \epsilon]$. \square

The required bit length L (the length of $\widehat{P}_i(m)$ values) is $\text{poly}(n, \epsilon^{-1}, \log \frac{1}{\Delta})$. So all numbers involved in our algorithms can be encoded in $\text{poly}(n, \epsilon^{-1}, \log C, \log \frac{1}{\Delta})$ bits, and the total number of arithmetic operations (additions, multiplications, comparisons) is also $\text{poly}(n, \epsilon^{-1}, \log C, \log \frac{1}{\Delta})$, we then conclude the overall bit complexity is $\text{poly}(n, \epsilon^{-1}, \log C, \log \frac{1}{\Delta})$.

Acknowledgments

We would like to thank Paul Tsui and the anonymous reviewer for their detailed comments, which helped us to improve the presentation significantly. This work was supported in part by the

National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, and the National Natural Science Foundation of China Grant 61202009, 61033001, 61061130540.

References

- [1] G. Bennett, Probability inequalities for the sum of independent random variables, *J. Amer. Statist. Assoc.* 57 (297) (1962) 33–45.
- [2] A. Bhalgat, A. Goel, S. Khanna, Improved approximation results for stochastic knapsack problems, in: ACM-SIAM Symposium on Discrete Algorithms, 2011.
- [3] H.E. Daniels, Tail probability approximations, in: *International Statistical Review/Revue Internationale de Statistique*, 1987, pp. 37–48.
- [4] M. Dyer, Approximate counting by dynamic programming, in: *ACM Symp. on Theory of Computing*, 2003, pp. 693–699.
- [5] S. Geetha, K. Nair, On stochastic spanning tree problem, *Networks* 23 (8) (1993) 675–679.
- [6] A. Goel, P. Indyk, Stochastic load balancing and related problems, in: *Annual Symp. on Foundations of Computer Science*, 1999, p. 579.
- [7] P. Gopalan, A. Klivans, R. Meka, D. Stefankovic, S. Vempala, E. Vigoda, An FPTAS for knapsack and related counting problems, in: *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, 2011, pp. 817–826.
- [8] V. Goyal, R. Ravi, Chance constrained knapsack problem with random item sizes, *Oper. Res. Lett.* (2009).
- [9] J. Kleinberg, Y. Rabani, É. Tardos, Allocating bandwidth for bursty connections, in: *ACM Symp. on Theory of Computing*, 1997, p. 673.
- [10] J. Li, A. Deshpande, Maximizing expected utility for stochastic combinatorial optimization problems, in: *Annual Symp. on Foundations of Computer Science*, 2011.
- [11] J. Li, W. Yuan, Stochastic combinatorial optimization via Poisson approximation, in: *ACM Symp. on Theory of Computing*, 2013.
- [12] R. Lugannani, S. Rice, Saddle point approximation for the distribution of the sum of independent random variables, *Adv. Appl. Probab.* (1980) 475–490.
- [13] N.B. Mehta, J. Wu, A.F. Molisch, J. Zhang, Approximating a sum of random variables with a lognormal, *IEEE Trans. Wirel. Commun.* 6 (7) (2007) 2690–2699.
- [14] M. Mitzenmacher, E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*, Cambridge University Press, 2005.
- [15] B. Morris, A. Sinclair, Random walks on truncated cubes and sampling 0–1 knapsack solutions, *SIAM J. Comput.* 34 (1) (2004) 195–226.
- [16] E. Nikolova, Approximation algorithms for reliable stochastic combinatorial optimization, in: *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2010, pp. 338–351.
- [17] E. Nikolova, J. Kelner, M. Brand, M. Mitzenmacher, Stochastic shortest paths via quasi-convex maximization, in: *European Symposium on Algorithms*, 2006, pp. 552–563.
- [18] V.V. Petrov, On the probabilities of large deviations for sums of independent random variables, *Theory Probab. Appl.* 10 (2) (1965) 287–298.
- [19] D. Štefankovic, S. Vempala, E. Vigoda, A deterministic polynomial-time approximation scheme for counting knapsack solutions, *SIAM J. Comput.* 41 (2) (2012) 356–366.