

THE INFORMATION-BASED COMPLEXITY OF DYNAMIC PROGRAMMING †

Chee-Seng Chow^{††}
John N. Tsitsiklis^{††}

Abstract

We provide tight lower bounds on the computational complexity of discrete time, stationary, infinite horizon, discounted stochastic control problems, for the case where the state space is continuous and the problem is to be solved approximately, within a specified accuracy. We study the dependence of the complexity on the desired accuracy and on the discount factor.

† Research supported by the NSF under Grant ECS-8552419, with matching funds from Bellcore and Dupont, and by the ARO under Grant DAAL03-86-K-0171.

†† Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139.

1. INTRODUCTION

This paper addresses issues related to the computational complexity of solving discrete-time stochastic control problems defined on a continuous state space. There has been a great deal of research on the computational aspects of stochastic control when the state space is finite ([B1,PT,HTW]). However, much less has been accomplished for the case of continuous state spaces. An explanation for this state of affairs could be that such problems are very demanding computationally, with most realistic problems lying beyond the capabilities of commercial computers. However, with advances in computer hardware and with the availability of new powerful architectures, it is to be expected that the numerical solution of continuous-state stochastic control problems will become much more common, hence the motivation for our work.

Let S be some subset of \mathfrak{R}^n which is the state space of a controlled stochastic process. A large class of discrete-time stochastic control problems boils down to the computation of a fixed point J^* of the nonlinear operator T (acting on a space of functions on the set S) defined by

$$(TJ)(x) = \inf_{u \in C} \left[g(x, u) + \alpha \int_S J(y)P(y|x, u) dy \right], \quad \forall x \in S. \quad (1.1)$$

Here, $C \subset \mathfrak{R}^m$ is the control space, $g(x, u)$ is the cost incurred if the current state is x and control u is applied, $\alpha \in (0, 1)$ is a discount factor, and $P(y|x, u)$ is a stochastic kernel that specifies the probability distribution of the next state y , when the current state is x and control u is applied. Then, $J^*(x)$ is interpreted as the value of the expected discounted cost, starting from state x , and provided that the control actions are chosen optimally (see Section 2).

A fixed point J^* of the operator T cannot be determined analytically except for a limited class of examples. On the other hand, an approximation to such a fixed point can be computed by suitably discretizing the state and control spaces, and then solving a finite-dimensional version of the problem. There has been some work on such discretization methods, with typical results demonstrating that as the discretization becomes finer, the resulting approximation of J^* becomes more and more accurate [Wh,B2,CT]. Upper bounds on the approximation error are also available. One of the consequences of the results to be derived in the present paper is that some of the earlier upper bounds are tight within a constant factor.

Once the original problem is discretized, there is a choice of numerical methods for solving the discrete problem [B1]. One particular choice is studied in a companion paper [CT], where it is shown that the total computational effort is closely related to the amount of work needed in discretizing the problem, when a suitable multigrid method is employed. Thus, the results in the present paper demonstrate that the algorithm in [CT] is close to optimal (and sometimes optimal) as far as its complexity is concerned.

In the special case where the control space U consists of a single element, the minimization in Eq. (1.1) is redundant, and the fixed point equation $J^* = TJ^*$ becomes a (linear) Fredholm equation of the second kind. Thus, the results in the present paper, in conjunction with the algorithms in [CT],

characterize the computational complexity of the approximate solution of Fredholm equations of the second kind. There has been some prior research on this subject [W]. Our work is different in a number of respects that are discussed in Section 4.

The paper is organized as follows. In Section 2, we introduce our assumptions, define the problem to be solved, and state its relation to stochastic control. We also outline the model of computation to be employed. Finally, we state some known upper bounds on the approximation error introduced by the discretization of the problem, and translate them to complexity bounds. In Section 3, which contains our main results, lower bounds are derived on the information-based complexity of our problem. Finally, Section 4 contains our conclusions and some discussion of related issues.

2. PROBLEM DEFINITION

Let $S = [0, 1]^n$ and $U = [0, 1]^m$. Let α be a scalar belonging to $(0, 1)$. Let $g : S \times U \mapsto \mathfrak{R}$ and $P : S \times S \times U \mapsto \mathfrak{R}$ be some functions. Let K be some positive constant.

Assumption 2.1:

- a) $0 \leq P(y | x, u), \quad \forall y, x \in S, \forall u \in U.$
- b) $|P(y | x, u) - P(y' | x', u')| \leq K \|(y, x, u) - (y', x', u')\|_\infty, \quad \forall y, y', x, x' \in S, \forall u, u' \in U.$
- c) $|g(x, u) - g(x', u')| \leq K \|(x, u) - (x', u')\|_\infty, \quad \forall x, x' \in S, \forall u, u' \in U.$
- d) $|g(x, u)| \leq K, \quad \forall x \in S, \forall u \in U.$

Assumption 2.2: $\int_S P(y | x, u) dy = 1, \quad \forall x \in S, \forall u \in C.$

According to Assumptions 2.1 and 2.2, for any fixed $x \in S$ and $u \in U$, the function $P(\cdot | x, u)$ is a probability density on the set S . Furthermore, P and g are Lipschitz continuous with Lipschitz constant K .

Let $\mathcal{C}(S)$ be the set of all continuous real-valued functions on the set S . We define the operator $T : \mathcal{C}(S) \mapsto \mathcal{C}(S)$ by letting

$$(TJ)(x) = \min_{u \in U} \left[g(x, u) + \alpha \int_S J(y) P(y | x, u) dy \right], \quad \forall J \in \mathcal{C}(S), \forall x \in S. \quad (2.1)$$

(The fact that T maps $\mathcal{C}(S)$ into itself is proved in [CT].) The space $\mathcal{C}(S)$ endowed with the norm $\|J\|_\infty = \max_{x \in S} |J(x)|$ is a Banach space. Furthermore, T has the contraction property

$$\|TJ - TJ'\|_\infty \leq \alpha \|J - J'\|_\infty, \quad \forall J, J' \in \mathcal{C}(S). \quad (2.2)$$

Since $\alpha \in (0, 1)$, T is a contraction operator, and therefore has a unique fixed point $J^* \in \mathcal{C}(S)$. The equation $TJ = J$, of which J^* is the unique solution, is known as *Bellman's equation* or as the *dynamic programming equation*. We are interested in the computational aspects of the approximate evaluation of J^* .

Stochastic Control Interpretation

Let Π be the set of all Borel measurable functions $\mu : S \mapsto U$. Let Π^∞ be the set of all sequences of elements of Π . An element $\pi = (\mu_0, \mu_1, \dots)$ of Π^∞ , called a *policy* (also known as *feedback law* or *control law*), is to be viewed as a prescription for choosing an action $\mu_t(x_t) \in U$ at time t , as a function of the current state x_t of a controlled stochastic process. More precisely, given a policy $\pi = (\mu_0, \mu_1, \dots)$, we define a (generally, nonstationary) Markov process $\{x_t^\pi \mid t = 0, 1, \dots\}$ on the state space S by letting $P(\cdot \mid x_t^\pi, \mu_t(x_t^\pi))$ be the probability density function of x_{t+1}^π , conditioned on x_t^π . The cost $J_\pi(x)$ associated to such a policy is defined (as a function of the initial state x) by

$$J_\pi(x) = E \left[\sum_{t=0}^{\infty} \alpha^t g(x_t^\pi, \mu_t(x_t^\pi)) \mid x_0^\pi = x \right]. \quad (2.3)$$

[Notice that the infinite sum is absolutely convergent and bounded by $K/(1-\alpha)$ because $\alpha \in (0, 1)$ and the function g is bounded by K .] For any $x \in S$, we define $\hat{J}(x)$ by letting

$$\hat{J}(x) = \inf_{\pi \in \Pi^\infty} J_\pi(x), \quad (2.4)$$

and this defines a function $\hat{J} : S \mapsto \mathfrak{R}$. This function, known as the *cost-to-go* or *value* function, represents the least possible cost as a function of the initial state of the controlled process. A policy $\pi \in \Pi^\infty$ is called *optimal* if $J_\pi(x) = \hat{J}(x)$ for all $x \in S$. The central result of dynamic programming¹ states that \hat{J} coincides with the fixed point J^* of the operator T . Furthermore, once J^* is available, it is straightforward to determine an optimal policy. [This is done as follows: Consider Eq. (2.1) with J replaced by J^* . For each $x \in S$, choose some u that attains the minimum in Eq. (2.1), and let $\mu_t(x) = u$ for each t .] This justifies our interest in the function J^* .

The case where the discount factor α approaches 1 from below is of substantial theoretical and practical interest. For example, as $\alpha \uparrow 1$, one obtains, in the limit, the solution to a certain “average cost problem” [B1]. Also, if one deals with a discounted continuous-time stochastic control problem and the time step is discretized, one obtains a discrete-time discounted problem in which the discount factor approaches 1 as the time discretization step is made smaller.

In practical stochastic control problems, the state and control spaces could be arbitrary subsets of \mathfrak{R}^n and \mathfrak{R}^m , respectively, and there could be state-dependent constraints on the allowed values of u . Such problems can only be harder than the special case studied here. Thus, the lower bounds to be derived in Section 3 apply more generally. A similar comment applies to the smoothness conditions on g and P that have been imposed in Assumption 2.1.

Further Assumptions

In the case of stochastic control problems defined on a finite state space it is known that the convergence of certain algorithms for computing J^* is much faster when the controlled process

1. This result requires certain technical assumptions. Assumption 2.1 turns out to be sufficient [BS], [CT].

satisfies certain *mixing conditions* [B1] and the same is true in our case as well [CT]. Our next assumption introduces a condition of this type.

Assumption 2.3: There exists a constant $\rho > 0$ such that

$$\int_S \min_{x \in S, u \in U} P(y|x, u) dy \geq \rho. \quad (2.5)$$

Intuitively, Assumption 2.3 states that no matter what the current state is and what control is applied, there are certain states in S (of positive Lebesgue measure) for which the probability density of being visited at the next time step is positive. This ensures that the effects of initial conditions are quickly washed out.

In an alternative class of stochastic control problems, at any given time there is a certain probability, depending on the current state and the control being applied, that the process is terminated and costs stop accruing. Such a formulation is captured by allowing $P(\cdot|x, u)$ to be a subprobability measure, as in the following assumption:

Assumption 2.4: $\int_S P(y|x, u) dy \leq 1, \quad \forall x \in S, \forall u \in U.$

Model of Computation

Our computational task is completely determined by the functions P and g , the discount factor α , and the desired accuracy ϵ . Accordingly, a tuple (P, g, α, ϵ) will be called an *instance*. We then define a *problem* as a class of instances. In our context, different problems will correspond to different choices of assumptions.

In order to talk meaningfully about the approximate computation of J^* , we need a suitable model of computation. We use a real-number model of computation [TWW,YN] in which a processor:

- a) Performs comparisons of real numbers or infinite precision arithmetic operations in unit time.
- b) Submits queries $(y, x, u) \in S \times S \times U$ to an “oracle” and receives as answers the values of $g(x, u)$ and $P(y|x, u)$. [We then say that the processor *samples* (y, x, u) .] Queries can be submitted at any time in the course of the computation and this allows the values of (y, x, u) in a query to depend on earlier computations or on the answers to earlier queries. (We will therefore be dealing with “adaptive” algorithms, in the sense of in [TWW].)

An algorithm in the above model of computation can be loosely defined as a program that determines the computations to be performed and the queries to be submitted. An algorithm is said to be *correct* (for a given problem) if for every instance (P, g, α, ϵ) of the problem, it outputs a piecewise constant function J such that $\|J^* - J\|_\infty \leq \epsilon$, in some prespecified format. A natural format for the representation of the output is as follows. The processor outputs a parameter h that signifies that the state space S has been partitioned into cubes of volume h^n , and then outputs the value of J on each one of these cubes, with the understanding that J is constant on each one of these cubes.

The complexity of an algorithm of above described type is defined as the sum of:

- a) The number of oracle queries;
- b) The number of arithmetic operations performed by the algorithm.

A very fruitful method for establishing lower bounds on the complexity of any algorithm consists of lower bounding the number of queries that have to be submitted for the desired accuracy to be attainable. The typical argument here is that if the number of queries is small then the available information on the problem being solved is insufficient. Accordingly, the least number of queries necessary is often called the *information-based* complexity of the problem.

Let the dimensions n , m of the state and control spaces be fixed and let us view the constants K and ρ of Assumptions 2.1–2.3 as absolute constants. We consider three different problems:

- a) Problem \mathcal{P}_{prob} that consists of all instances that satisfy Assumptions 2.1 and 2.2. [In particular, $P(\cdot | x, u)$ is a probability measure for all (x, u) .]
- b) Problem \mathcal{P}_{mix} that consists of all instances that satisfy Assumptions 2.1–2.3. (That is, a mixing condition is also in effect.)
- c) Problem \mathcal{P}_{sub} that consists of all instances that satisfy Assumptions 2.1, 2.3–2.4. [That is, the mixing condition is still in effect, but $P(\cdot | x, u)$ is only a subprobability measure.]

Let us fix some ϵ and α . We define $C_{prob}(\alpha, \epsilon)$ as the minimum (over all correct algorithms for the problem \mathcal{P}_{prob}) of the number of queries, in the worst case over all instances (P, g, α, ϵ) of \mathcal{P}_{prob} . The quantities $C_{mix}(\alpha, \epsilon)$ and $C_{sub}(\alpha, \epsilon)$ are defined similarly, by replacing problem \mathcal{P}_{prob} by \mathcal{P}_{mix} and \mathcal{P}_{sub} , respectively.

The following upper bounds, together with discretization procedures that stay within these bounds, can be found in [Wh] and [CT]:

$$C_{prob}(\alpha, \epsilon) = O\left(\frac{1}{((1-\alpha)^2\epsilon)^{2n+m}}\right). \quad (2.6)$$

$$C_{mix}(\alpha, \epsilon) = O\left(\frac{1}{((1-\alpha)\epsilon)^{2n+m}}\right). \quad (2.7)$$

$$C_{sub}(\alpha, \epsilon) = O\left(\frac{1}{((1-\alpha)^2\epsilon)^{2n+m}}\right). \quad (2.8)$$

We have used the $O(\cdot)$ notation which should be interpreted as follows. Let f and h be functions from $(0, 1) \times (0, \infty)$ into $[0, \infty)$. We write $f(\alpha, \epsilon) = O(h(\alpha, \epsilon))$ if there exist constants $c > 0$, $\epsilon_0 > 0$, and $\alpha_0 \in (0, 1)$ such that $f(\alpha, \epsilon) \leq ch(\alpha, \epsilon)$ for all $\epsilon \in (0, \epsilon_0)$ and all $\alpha \in (\alpha_0, 1)$. (These constants are allowed to depend on the absolute constants n , m , K , and ρ .) Later, we will also use the notation $f(\alpha, \epsilon) = \Omega(h(\alpha, \epsilon))$ which is equivalent to $h(\alpha, \epsilon) = \Omega(f(\alpha, \epsilon))$.

3. LOWER BOUNDS

In this section, we prove that the upper bounds of Eqs. (2.6)–(2.8) are tight, by establishing the corresponding lower bounds. Our results rest on an “adversary” argument that is very common in the study of information-based complexity [TWW]. The outline of the argument is as follows. Suppose that a certain algorithm makes at most A queries. We consider a particular instance (P, g, α, ϵ) and we let X be the set of triples (y, x, u) sampled by the algorithm when presented with that instance. We then construct an alternative instance $(\hat{P}, \hat{g}, \alpha, \epsilon)$ such that $P(y|x, u) = \hat{P}(y|x, u)$ and $g(x, u) = \hat{g}(x, u)$ for all $(y, x, u) \in X$. The algorithm has no means of distinguishing between the two instances and must produce the same output J in both cases. Let J^* and \hat{J}^* be the optimal cost functions for the two problems. If we can manage so that $\|J^* - \hat{J}^*\|_\infty > 2\epsilon$, then at least one of the inequalities $\|J - J^*\|_\infty > \epsilon$ and $\|J - \hat{J}^*\|_\infty > \epsilon$ must hold. It follows that the algorithm cannot succeed for all problem instances, and therefore the information-based complexity of the problem is larger than the cardinality of A .

Theorem 3.1: (Lower bound under Assumptions 2.1–2.3) For any $K > 0$, $\rho \in (0, 1)$, m , and n , we have

$$C_{mix}(\alpha, \epsilon) = \Omega\left(\frac{1}{((1-\alpha)\epsilon)^{2n+m}}\right).$$

Proof: We only prove the result for the case $K = 1$ and $\rho = 1/2$. The proof for the general case is identical except for a minor modification discussed at the end of the proof. Let us also fix the dimensions m, n of the problem. Throughout the proof, an *absolute constant* will stand for a constant that can only depend on m, n , but not on any other parameters.

We fix some $\epsilon > 0$ and some $\alpha \in (1/2, 1)$. Let us consider some algorithm that is correct for the problem \mathcal{P}_{mix} and suppose that the number of queries is at most A for every instance with those particular values of α and ϵ . We will derive a lower bound on A .

We choose a positive scalar δ so that $1/\delta$ is an integer multiple of 16 and such that

$$\frac{1}{\delta_0} \leq \frac{1}{\delta} \leq \frac{1}{\delta_0} + 16, \quad (3.1)$$

where δ_0 satisfies

$$A = \frac{1}{4\delta_0^{2n+m}}. \quad (3.2)$$

We partition the set $S \times S \times U$ into cubic cells of volume δ^{2n+m} . (In particular, there will be $1/\delta^{2n+m}$ cells.) This is done by first specifying the “centers” of the cells. Let \tilde{S} be the set of all $x = (x_1, \dots, x_n) \in S$ such that each component x_i is of the form $x_i = (t + (1/2))\delta$, where t is a nonnegative integer smaller than $1/\delta$. Similarly, we let \tilde{U} be the set of all $u = (u_1, \dots, u_m) \in U$ such that each u_i is of the form $u_i = (t + (1/2))\delta$, where again t is a nonnegative integer smaller than $1/\delta$. For any $(\tilde{y}, \tilde{x}, \tilde{u}) \in \tilde{S} \times \tilde{S} \times \tilde{U}$, we define the cell $C_{\tilde{y}, \tilde{x}, \tilde{u}}$ by letting

$$C_{\tilde{y}, \tilde{x}, \tilde{u}} = \left\{ (y, x, u) \in S \times S \times U \mid \|(y, x, u) - (\tilde{y}, \tilde{x}, \tilde{u})\|_\infty < \frac{\delta}{2} \right\}.$$

Clearly, the cardinality of \tilde{S} and \tilde{U} is $1/\delta^n$ and $1/\delta^m$, respectively. It follows that there is a total of $1/\delta^{2n+m}$ cells. Notice that distinct cells are disjoint.

For any $(\tilde{y}, \tilde{x}, \tilde{u}) \in \tilde{S} \times \tilde{S} \times \tilde{U}$, we define a function $E_{\tilde{y}, \tilde{x}, \tilde{u}} : S \times S \times U \mapsto \mathfrak{R}$, by letting

$$E_{\tilde{y}, \tilde{x}, \tilde{u}}(\mathbf{y} | \mathbf{x}, \mathbf{u}) = 0, \quad \text{if } (\mathbf{y}, \mathbf{x}, \mathbf{u}) \notin C_{\tilde{y}, \tilde{x}, \tilde{u}},$$

$$E_{\tilde{y}, \tilde{x}, \tilde{u}}(\mathbf{y} | \mathbf{x}, \mathbf{u}) = \frac{\delta}{2} - \|(\mathbf{y}, \mathbf{x}, \mathbf{u}) - (\tilde{y}, \tilde{x}, \tilde{u})\|_\infty, \quad \text{if } (\mathbf{y}, \mathbf{x}, \mathbf{u}) \in C_{\tilde{y}, \tilde{x}, \tilde{u}}. \quad (3.3)$$

Thus, $E_{\tilde{y}, \tilde{x}, \tilde{u}}$ is just a ‘‘pyramid’’ of height $\delta/2$ whose base is the cell $C_{\tilde{y}, \tilde{x}, \tilde{u}}$. The triangle inequality applied to the norm $\|\cdot\|_\infty$ shows that

$$|E_{\tilde{y}, \tilde{x}, \tilde{u}}(\mathbf{y} | \mathbf{x}, \mathbf{u}) - E_{\tilde{y}, \tilde{x}, \tilde{u}}(\mathbf{y}' | \mathbf{x}', \mathbf{u}')| \leq \|(\mathbf{y}, \mathbf{x}, \mathbf{u}) - (\mathbf{y}', \mathbf{x}', \mathbf{u}')\|_\infty, \quad \forall (\mathbf{y}, \mathbf{x}, \mathbf{u}), (\mathbf{y}', \mathbf{x}', \mathbf{u}') \in C_{\tilde{y}, \tilde{x}, \tilde{u}}.$$

Thus, $E_{\tilde{y}, \tilde{x}, \tilde{u}}$ satisfies the Lipschitz continuity Assumption 2.1, with Lipschitz constant $K = 1$, on the set $C_{\tilde{y}, \tilde{x}, \tilde{u}}$. The function $E_{\tilde{y}, \tilde{x}, \tilde{u}}$ is continuous at the boundary of $C_{\tilde{y}, \tilde{x}, \tilde{u}}$ and is zero outside $C_{\tilde{y}, \tilde{x}, \tilde{u}}$. Thus, $E_{\tilde{y}, \tilde{x}, \tilde{u}}$ is obtained by piecing together in a continuous manner a Lipschitz continuous function and a constant function. It follows that $E_{\tilde{y}, \tilde{x}, \tilde{u}}$ is Lipschitz continuous on the set $S \times S \times U$, with Lipschitz constant 1.

We define an instance (P, g, α, ϵ) by letting

$$g(\mathbf{x}, \mathbf{u}) = x_1, \quad \forall (\mathbf{x}, \mathbf{u}) \in S \times U, \quad (3.4)$$

and

$$P(\mathbf{y} | \mathbf{x}, \mathbf{u}) = 1, \quad \forall (\mathbf{y}, \mathbf{x}, \mathbf{u}) \in S \times S \times U. \quad (3.5)$$

It is easily seen that this instance satisfies Assumptions 2.1–2.3 with $K = 1$ and $\rho = 1/2$.

Bellman’s equation reads

$$J(\mathbf{x}) = x_1 + \alpha \int_S J(\mathbf{y}) d\mathbf{y}. \quad (3.6)$$

A simple calculation shows that the function J^* defined by

$$J^*(\mathbf{x}) = x_1 + \frac{\alpha}{2(1 - \alpha)} \quad (3.7)$$

is a solution of (3.6) and according to the discussion of Section 2, it is the unique fixed point of T .

Let X be the set of points $(\mathbf{y}, \mathbf{x}, \mathbf{u})$ sampled by the particular algorithm we are considering, when it is faced with the instance (P, g, α, ϵ) . In particular, the cardinality of X is at most A . Using the definition of δ [cf. Eqs. (3.1)–(3.2)], the cardinality of X is at most $1/(4\delta^{2n+m})$.

We say that a cell $C_{\tilde{y}, \tilde{x}, \tilde{u}}$ is *sampled* if the intersection of X and $C_{\tilde{y}, \tilde{x}, \tilde{u}}$ is nonempty. Otherwise, we say that $C_{\tilde{y}, \tilde{x}, \tilde{u}}$ is *unsampled*. We say that some $(\tilde{x}, \tilde{u}) \in \tilde{S} \times \tilde{U}$ is *well-sampled* if there exist at least $1/(2\delta^n)$ elements \tilde{y} of \tilde{S} for which the cell $C_{\tilde{y}, \tilde{x}, \tilde{u}}$ is sampled. Otherwise, we say that (\tilde{x}, \tilde{u}) is *badly sampled*. Since the total number of samples is bounded by $1/(4\delta^{2n+m})$, there exist at most

$1/(2\delta^{n+m})$ well-sampled elements $(\tilde{x}, \tilde{u}) \in \tilde{S} \times \tilde{U}$. Therefore, there are at least $1/(2\delta^{n+m})$ badly sampled (\tilde{x}, \tilde{u}) . For each $\tilde{x} \in \tilde{S}$ there are at most $1/\delta^m$ possible choices of \tilde{u} such that (\tilde{x}, \tilde{u}) is badly sampled. This shows that there exists a set $\tilde{S}_{\text{BAD}} \subset \tilde{S}$ of cardinality $1/(2\delta^n)$ such that for each $\tilde{x} \in \tilde{S}_{\text{BAD}}$ there exists some $\tilde{\mu}(\tilde{x}) \in \tilde{U}$ for which $(\tilde{x}, \tilde{\mu}(\tilde{x}))$ is badly sampled.

We will now construct a second instance. The cost function g is left unchanged [cf. Eq. (3.4)], but we modify the probability density on some of the unsampled cells. This is done as follows. Let us fix some $\tilde{x} \in \tilde{S}_{\text{BAD}}$. By the definition of \tilde{S}_{BAD} and $\tilde{\mu}(\tilde{x})$, if we keep \tilde{x} fixed and vary \tilde{y} , we find at least $1/(2\delta^n)$ unsampled cells of the form $C_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}$. We sort these unsampled cells in order of increasing \tilde{y}_1 and we let c be the median value of \tilde{y}_1 . We refer to those cells for which $\tilde{y}_1 \leq c$ (respectively, $\tilde{y}_1 \geq c$) as *low* (respectively, *high*) cells. Let $\underline{c} = c - (1/16)$ and $\bar{c} = c + (1/16)$. We discard all unsampled cells $C_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}$ for which $\underline{c} < \tilde{y}_1 < \bar{c}$. Thus, the number of discarded cells is bounded by $1/(8\delta^n)$. Since we started with at least $1/(4\delta^n)$ low unsampled cells, we are left with at least $1/(8\delta^n)$ such cells. By discarding some more low unsampled cells (if needed), we can assume that we are left with exactly $1/(8\delta^n)$ unsampled low cells. By a similar argument, we can also assume that we are left with exactly $1/(8\delta^n)$ unsampled high cells. Let $Q^L(\tilde{x})$ [respectively, $Q^H(\tilde{x})$] be the set of all $\tilde{y} \in \tilde{S}$ such that $C_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}$ is a low (respectively, high) unsampled cell that has not been discarded. This procedure is carried out for each $\tilde{x} \in \tilde{S}_{\text{BAD}}$.

We define

$$\hat{P}(y | x, u) = P(y | x, u) + E(y | x, u) = 1 + E(y | x, u), \quad (3.8)$$

where

$$E(y | x, u) = \sum_{\tilde{x} \in \tilde{S}_{\text{BAD}}, \tilde{y} \in Q^L(\tilde{x})} E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, u) - \sum_{\tilde{x} \in \tilde{S}_{\text{BAD}}, \tilde{y} \in Q^H(\tilde{x})} E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, u). \quad (3.9)$$

In words, we add a pyramid at each low unsampled cell and we subtract a pyramid at each high unsampled cell. This has the effect of shifting the transition probability distribution closer to the origin, with a consequent decrease in the cost incurred after a transition.

We verify that our perturbed instance $(\hat{P}, g, \alpha, \epsilon)$ satisfies the required assumptions. Since each pyramid is Lipschitz continuous with Lipschitz constant 1, and since distinct pyramids are supported on distinct cells, it follows that Assumption 2.1(b) is satisfied with $K = 1$. Furthermore, for each (x, u) , the number of added pyramids is equal to the number of subtracted pyramids. For this reason, $\int_S E(y | x, u) dy = 0$ and \hat{P} satisfies Assumption 2.2. Finally, the height of each pyramid is $\delta/2$. Since $\delta \leq 1$, we have $\hat{P}(y | x, u) \geq 1 - (\delta/2) \geq 1/2$. This shows that Assumption 2.2 and Assumption 2.3 (with $\rho = 1/2$) are satisfied.

Our next task is to estimate the optimal cost function \hat{J}^* corresponding to the perturbed instance $(\hat{P}, g, \alpha, \epsilon)$. Let

$$B = \left\{ x \in S \mid \exists \tilde{x} \in \tilde{S}_{\text{BAD}} \text{ such that } \|x - \tilde{x}\|_\infty \leq \frac{\delta}{4} \right\}.$$

For any $x \in B$, we let $\mu(x) = \tilde{\mu}(\tilde{x})$ where \tilde{x} is the element of \tilde{S}_{BAD} for which $\|x - \tilde{x}\|_\infty \leq \delta/4$. For any $x \notin B$, we let $\mu(x) = 0$. We now consider the quantity

$$e(x) = \int_S g(y) E(y | x, \mu(x)) dy \quad (3.10)$$

that can be interpreted as the effect of the perturbation on the expected cost after the first transition, when the control is chosen according to the function μ .

Lemma 3.1: For each $x \in S$, we have $e(x) \leq 0$. Furthermore, there exists a positive absolute constant k such that $e(x) \leq -k\delta$, for all $x \in B$.

Proof: Using Eqs. (3.9) and (3.10) and the definition of g , we have

$$e(x) = \sum_{\tilde{x} \in \tilde{S}_{\text{BAD}}, \tilde{y} \in Q^L(\tilde{x})} \int_S y_1 E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, \mu(x)) dy - \sum_{\tilde{x} \in \tilde{S}_{\text{BAD}}, \tilde{y} \in Q^H(\tilde{x})} \int_S y_1 E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, \mu(x)) dy. \quad (3.11)$$

For any $x \notin B$, we have $\mu(x) = 0$ which implies that $E(y | x, \mu(x)) = 0$ and $e(x) = 0$. Let us now fix some $x \in B$ and let \tilde{x} be the corresponding element of \tilde{S}_{BAD} . Then, Eq. (3.11) becomes

$$e(x) = \sum_{\tilde{y} \in Q^L(\tilde{x})} \int_S y_1 E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, \mu(x)) dy - \sum_{\tilde{y} \in Q^H(\tilde{x})} \int_S y_1 E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, \mu(x)) dy. \quad (3.12)$$

Let us consider the summand corresponding to a particular $\tilde{y} \in Q^L(\tilde{x})$. We only need to carry out the integration on the set $Y(\tilde{y}) = \{y \in S \mid \|y - \tilde{y}\|_\infty \leq \delta/2\}$ (instead of the entire set S) because $E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, \mu(x))$ vanishes when $y \notin Y(\tilde{y})$. For $y \in Y(\tilde{y})$, we have $y_1 \leq \tilde{y}_1 + \delta/2 \leq \underline{c} + \delta/2$. We now use the definition of the function $E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, \mu(x))$ [cf. Eq. (3.3)], together with the property $\mu(x) = \tilde{\mu}(\tilde{x})$, to conclude that

$$\int_S y_1 E_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}(y | x, \mu(x)) dy \leq \left(\underline{c} + \frac{\delta}{2}\right) I(x),$$

where

$$I(x) = \int_{Y(\tilde{y})} \left(\frac{\delta}{2} - \max\{\|x - \tilde{x}\|_\infty, \|y - \tilde{y}\|_\infty\}\right) dy. \quad (3.13)$$

It is clear that the value of $I(x)$ is independent of the choice of \tilde{y} , which justifies our notation. By a symmetrical argument, each one of the summands corresponding to $\tilde{y} \in Q^H(\tilde{x})$ is bounded below by $(\bar{c} - \delta/2)I(x)$. Since each one of the sets $Q^H(\tilde{x})$, $Q^L(\tilde{x})$ has cardinality $1/8\delta^n$, it follows from Eq. (3.11) that

$$e(x) \leq -(\bar{c} - \underline{c} - \delta)I(x) \frac{1}{8\delta^n} \leq -I(x) \frac{1}{64\delta^n}, \quad \forall x \in B, \quad (3.14)$$

where the last inequality follows because $\bar{c} - \underline{c} = 1/4$ (by construction) and $\delta \leq 1/16$ (by definition). We now bound $I(x)$ for $x \in B$. We have $\|x - \tilde{x}\|_\infty \leq \delta/4$, the integrand is always nonnegative, and is at least $\delta/4$ for every y belonging to the set $\{y \in S \mid \|y - \tilde{y}\|_\infty \leq \delta/4\}$. Therefore, for $x \in B$,

$I(x)$ is bounded below by $\delta/4$ times the volume of the set $\{\mathbf{y} \in S \mid \|\mathbf{y} - \tilde{\mathbf{y}}\|_\infty \leq \delta/4\}$. This set is an n -dimensional cube, whose edges have length $\delta/2$. Thus, we obtain

$$I(x) \geq \frac{\delta \delta^n}{4 \cdot 2^n}, \quad \forall x \in B.$$

Combining with Eq. (3.14), we obtain

$$e(x) \leq -\frac{\delta}{4 \cdot 64 \cdot 2^n}, \quad \forall x \in B,$$

which proves the desired result. **q.e.d.**

Lemma 3.2: There exists a positive absolute constant h such that

$$\int_B \hat{P}(\mathbf{y} \mid x, \mathbf{u}) d\mathbf{y} \geq h, \quad \forall x \in S, \forall \mathbf{u} \in U.$$

Proof: The function \hat{P} is bounded below by $1/2$, as shown earlier. Thus, it suffices to show that the volume of B is bounded below by some absolute constant. Notice that B consists of $1/(2\delta^n)$ cubes of volume $(\delta/2)^n$, and the result follows. **q.e.d.**

Let \hat{T} be the operator defined by Eq. (2.1) but with P replaced by \hat{P} . We have

$$\begin{aligned} \hat{T}J^*(x) &= g(x) + \alpha \min_{\mathbf{u} \in U} \int_S J^*(\mathbf{y}) \hat{P}(\mathbf{y} \mid x, \mathbf{u}) d\mathbf{y} \\ &\leq g(x) + \alpha \int_S J^*(\mathbf{y}) \hat{P}(\mathbf{y} \mid x, \mu(x)) d\mathbf{y} \\ &= x_1 + \alpha \int_S J^*(\mathbf{y}) d\mathbf{y} + \alpha \int_S J^*(\mathbf{y}) E(\mathbf{y} \mid x, \mu(x)) d\mathbf{y} \\ &= J^*(x) + \alpha \int_S J^*(\mathbf{y}) E(\mathbf{y} \mid x, \mu(x)) d\mathbf{y} \\ &= J^*(x) + \alpha \int_S \mathbf{y}_1 E(\mathbf{y} \mid x, \mu(x)) d\mathbf{y} + \frac{\alpha^2}{2(1-\alpha)} \int_S E(\mathbf{y} \mid x, \mu(x)) d\mathbf{y} \\ &= J^*(x) + \alpha e(x), \end{aligned} \tag{3.15}$$

where we have used the fact that J^* satisfies Eqs. (3.6) and (3.7), the definition of $e(x)$ [cf. Eq. (3.10)], and the fact that $\int_S E(\mathbf{y} \mid x, \mu(x)) d\mathbf{y} = 0$. It follows that

$$\hat{T}J^*(x) \leq J^*(x), \quad \forall x \in S, \tag{3.16}$$

$$\hat{T}J^*(x) \leq J^*(x) - \alpha k \delta, \quad \forall x \in B, \tag{3.17}$$

where k is the constant of Lemma 3.1. Let \hat{T}^t be the composition of t copies of \hat{T} and let $\bar{B} = \{x \in$

$S \setminus \{x \notin B\}$ be the complement of B . We have

$$\begin{aligned}
\hat{T}^2 J^*(x) &= g(x) + \alpha \min_{u \in U} \int_S \hat{T} J^*(y) \hat{P}(y|x, u) dy \\
&\leq g(x) + \alpha \int_S \hat{T} J^*(y) \hat{P}(y|x, \mu(x)) dy \\
&= g(x) + \alpha \int_B \hat{T} J^*(y) \hat{P}(y|x, \mu(x)) dy + \alpha \int_{\bar{B}} \hat{T} J^*(y) \hat{P}(y|x, \mu(x)) dy \\
&\leq g(x) + \alpha \int_B (J^*(y) - \alpha k \delta) \hat{P}(y|x, \mu(x)) dy + \alpha \int_{\bar{B}} J^*(y) \hat{P}(y|x, \mu(x)) dy \\
&= J^*(x) + \alpha e(x) - \alpha^2 k \delta \int_B \hat{P}(y|x, \mu(x)) dy \\
&\leq J^*(x) - \alpha^2 k \delta h, \quad \forall x \in S.
\end{aligned} \tag{3.18}$$

[We have used here the equality between the second and the last line of Eq. (3.15), as well as Lemma 3.2.] It is well known [B1] (and easy to verify) that for any real constant d , we have $\hat{T}(J+d) = \alpha d + \hat{T}J$. (Here the notation $J+d$ should be interpreted as the function which is equal to the sum of J with a function on S that is identically equal to d .) Using this property and Eq. (3.18), we obtain

$$\hat{T}^3 J^*(x) \leq \hat{T}^2 J^*(x) - \alpha^3 k \delta h \leq J^*(x) - \alpha^2 k \delta h - \alpha^3 k \delta h, \quad \forall x \in S.$$

We continue inductively, to obtain

$$\hat{T}^t J^*(x) \leq J^*(x) - (1 + \alpha + \alpha^2 + \dots + \alpha^{t-2}) \alpha^2 k \delta h, \quad t = 2, 3, \dots, \forall x \in S. \tag{3.19}$$

Taking the limit as $t \rightarrow \infty$, $\hat{T}^t J^*$ converges to the optimal cost function \hat{J}^* of the perturbed instance, and Eq. (3.19) implies that

$$\hat{J}^*(x) \leq J^*(x) - \frac{\alpha^2}{1 - \alpha} \delta k h, \quad \forall x \in S. \tag{3.20}$$

Notice that the perturbed instance coincides with the original one at all points sampled by the algorithm. For this reason, the algorithm will perform the same arithmetic operations and will return the same answer for both instances. That answer must be an ϵ -approximation of both J^* and \hat{J}^* . It follows that $\|J^* - \hat{J}^*\|_\infty \leq 2\epsilon$. Therefore,

$$\frac{\alpha^2}{1 - \alpha} \delta k h \leq 2\epsilon.$$

Since $\alpha \geq 1/2$, we obtain

$$\delta \leq d(1 - \alpha)\epsilon, \tag{3.21}$$

where d is some absolute constant.




Figure 1. The functions F_1 , F_2 , and G . The maximum value a of F_1 and F_2 is chosen so that $\int_0^1 F_1(x) dx = \int_0^1 F_2(x) dx = 1$. In particular, $3 \leq a \leq 9/2$.

For $\epsilon \leq 1/(32d)$, we obtain $\delta \leq 1/32$ or $1/(2\delta) \geq 16$. Thus, using Eq. (3.1), we have $1/\delta_0 \geq (1/\delta) - 16 \geq (1/\delta) - (1/2\delta) = 1/(2\delta)$, and Eq. (3.2) yields

$$A = \frac{1}{4\delta_0^{2n+m}} \geq \frac{1}{4(2\delta)^{2n+m}} \geq \frac{1}{4(2d(1-\alpha)\epsilon)^{2n+m}} = \Omega\left(\frac{1}{((1-\alpha)\epsilon)^{2n+m}}\right).$$

When the theorem is proved for general values of K and ρ , it is sufficient to multiply the pyramidal functions of Eq. (3.3) by a factor of $\min\{K, 1-\rho\}$. It is then easily seen that the perturbed problem satisfies Assumptions 2.1 and 2.3 for the given values of K and ρ and the proof goes through verbatim, except that certain absolute constants are modified. **Q.E.D.**

In our next result, the mixing condition (Assumption 2.3) is removed. It will be seen that this allows us to obtain a larger lower bound.

Theorem 3.2: (Lower bound under Assumptions 2.1–2.2) For every m, n , there exists some K such that

$$C_{prob}(\epsilon, \alpha) = \Omega\left(\frac{1}{((1-\alpha)^2\epsilon)^{2n+m}}\right).$$

Proof: The structure of the proof is similar to the preceding one. We fix n, m , and some K that will depend on n in a way to be determined later. An absolute constant is again a constant that depends only on m and n .

We fix some $\epsilon > 0$ and some $\alpha \in (1/2, 1)$. We consider an algorithm that is correct for the problem \mathcal{P}_{prob} (for the given values of m, n, K) and suppose that the number of queries is at most A for every instance with those particular values of α and ϵ .

We choose a positive scalar δ so that $1/\delta$ is an integer multiple of 9 and such that

$$\frac{1}{\delta_0} \leq \frac{1}{\delta} \leq \frac{1}{\delta_0} + 9, \quad (3.22)$$

where δ_0 satisfies

$$A = \frac{1}{5} \left(\frac{2}{9}\right)^{2n} \frac{1}{\delta_0^{2n+m}}. \quad (3.23)$$

We partition $S \times S \times U$ into cubic cells of volume $1/\delta^{2n+m}$ exactly as in the proof of Theorem 3.1 and we use the same notations $\tilde{S}, \tilde{U}, C_{\tilde{y}, \tilde{x}, \tilde{u}}$, and $E_{\tilde{y}, \tilde{x}, \tilde{u}}$.

We define the first instance to be considered. Let $F_1, F_2, G : [0, 1] \mapsto \mathfrak{R}$ be the functions shown in Fig. 1. We define a function $H : [0, 1] \times [0, 1] \mapsto \mathfrak{R}$ by letting

$$H(y | x) = F_1(y)G(x) + F_2(y)(1 - G(x)), \quad \forall x, y \in [0, 1].$$

We finally let

$$P(y | x, u) = \prod_{i=1}^n H(y_i | x_i), \quad \forall (y, x, u) \in S \times S \times U, \quad (3.24)$$

where x_i and y_i is the i th component of x and y , respectively. As for the cost function g , we only assume that $g(x, u) = 1$ for all $x \in [0, 1/3]^n$ and $u \in U$, and that $g(x, u) = 0$ for all $x \in [2/3, 1]^n$ and $u \in U$.

We verify that Assumptions 2.1–2.2 are satisfied. The function P is certainly nonnegative. Furthermore, F_1 and F_2 integrate to 1. Consequently, $\int_{[0,1]} H(y | x) dy = 1$, for all $x \in [0, 1]$. Thus, for any x, u , $P(\cdot | x, u)$ is a product of probability measures [cf. Eq. (3.24)] and is itself a probability measure. Notice that F_1, F_2 and G are Lipschitz continuous. It follows that P is also Lipschitz continuous with Lipschitz constant K , provided that the absolute constant K is taken large enough. Concerning the function g , we have not specified it in detail, but it is easily seen that there exist Lipschitz continuous functions satisfying the requirements we have imposed on g .

Notice that the Markov chain corresponding to P has the property that if the current state is in the set $[0, 1/3]^n$ then the state stays forever in that set. The same property holds for the set $[2/3, 1]^n$.

We now estimate $J^*(x)$ when $x \in [0, 1/3]^n$. While we could argue directly in terms of Bellman's equation, the argument is much more transparent if we use the interpretation of $J^*(x)$ as the optimal cost expressed as a function of the initial state. Starting with some initial state in $[0, 1/3]^n$, the state never exits that set. Furthermore, $g(x) = 1$ for every $x \in [0, 1/3]^n$. This implies that

$$J^*(x) = \sum_{t=1}^{\infty} \alpha^t = \frac{1}{1 - \alpha}, \quad \forall x \in [0, 1/3]^n. \quad (3.25)$$

Lemma 3.3: There exists a set $\tilde{S}_{\text{BAD}} \subset [0, 2/9]^n \cap \tilde{S}$ of cardinality $(2/9)^n / (2\delta^n)$ with the following property: for every $\tilde{x} \in \tilde{S}_{\text{BAD}}$ there exists some $\tilde{\mu}(\tilde{x}) \in \tilde{U}$ and two sets $Q^L(\tilde{x}) \subset [0, 2/9]^n \cap \tilde{S}$, $Q^H(\tilde{x}) \subset [7/9, 1]^n \cap \tilde{S}$, of cardinality $(2/9)^n / (2\delta^n)$, such that the cell $C_{\tilde{y}, \tilde{x}, \tilde{\mu}(\tilde{x})}$ is unsampled for every $\tilde{y} \in Q^L(\tilde{x}) \cup Q^H(\tilde{x})$.

Proof: Let \tilde{S}_{GOOD} be the set of all $\tilde{x} \in [0, 2/9]^n \cap \tilde{S}$ that do not have the desired property. Since the cardinality of $[0, 2/9]^n \cap \tilde{S}$ is $(2/9)^n / \delta^n$ it is sufficient to show that \tilde{S}_{GOOD} has cardinality less than or equal to $(2/9)^n / (2\delta^n)$. We suppose the contrary, and we will obtain a contradiction.

Fix some $\tilde{x} \in \tilde{S}_{\text{GOOD}}$. Then, for every $\tilde{u} \in \tilde{U}$ we can find at least $(2/9)^n / (2\delta^n)$ values of $\tilde{y} \in \tilde{S}$ such that the cell $C_{\tilde{y}, \tilde{x}, \tilde{u}}$ is sampled. This shows that the total number of sampled cells is at least

$(2/9)^{2n}/(4\delta^{2n+m})$. Using Eqs. (3.22) and (3.23), this implies that the number of sampled cells is more than A , a contradiction. **q.e.d.**

We now construct a perturbed instance. The cost function g is left unchanged. We define

$$\hat{P}(\mathbf{y} \mid \mathbf{x}, \mathbf{u}) = P(\mathbf{y} \mid \mathbf{x}, \mathbf{u}) + E(\mathbf{y} \mid \mathbf{x}, \mathbf{u}),$$

where

$$E(\mathbf{y} \mid \mathbf{x}, \mathbf{u}) = \sum_{\tilde{\mathbf{x}} \in \tilde{S}_{\text{BAD}}, \tilde{\mathbf{y}} \in Q^H(\tilde{\mathbf{x}})} E_{\tilde{\mathbf{y}}, \tilde{\mathbf{x}}, \tilde{\mu}(\tilde{\mathbf{x}})}(\mathbf{y} \mid \mathbf{x}, \mathbf{u}) - \sum_{\tilde{\mathbf{x}} \in \tilde{S}_{\text{BAD}}, \tilde{\mathbf{y}} \in Q^L(\tilde{\mathbf{x}})} E_{\tilde{\mathbf{y}}, \tilde{\mathbf{x}}, \tilde{\mu}(\tilde{\mathbf{x}})}(\mathbf{y} \mid \mathbf{x}, \mathbf{u}). \quad (3.26)$$

In effect, we are giving positive probability to certain transitions from the set $[0, 2/9]^n$ to the set $[7/9, 1]^n$. On the other hand, the property that the state can never exit from the set $[7/9, 1]^n$ is retained. The Lipschitz continuity of E and P implies that \hat{P} is Lipschitz continuous. Also $\hat{P}(\cdot \mid \mathbf{x}, \mathbf{u})$ is nonnegative and integrates to 1, for reasons similar to those in the proof of Theorem 3.1. Thus, Assumptions 2.1 and 2.2 are satisfied.

Let

$$B = \left\{ \mathbf{x} \in [0, 1/3]^n \mid \exists \tilde{\mathbf{x}} \in \tilde{S}_{\text{BAD}} \text{ such that } \|\mathbf{x} - \tilde{\mathbf{x}}\|_\infty \leq \frac{\delta}{4} \right\}.$$

Lemma 3.4: For every $\mathbf{x} \in [0, 1/3]^n$, we have

$$\int_B \hat{P}(\mathbf{y} \mid \mathbf{x}, \mathbf{u}) \, d\mathbf{y} \geq h,$$

where h is a positive absolute constant.

Proof: Fix some $\mathbf{x} \in [0, 1/3]^n$. Notice that $P(\mathbf{y} \mid \mathbf{x}, \mathbf{u}) = \prod_{i=1}^n F_1(y_i) \geq 3^n \geq 3$ for all $\mathbf{y} \in [0, 2/9]^n$. Since $|E(\mathbf{y} \mid \mathbf{x}, \mathbf{u})| \leq \delta/2 \leq 1$, we conclude that $\hat{P}(\mathbf{y} \mid \mathbf{x}, \mathbf{u}) \geq 2$, for all $\mathbf{y} \in [0, 2/9]^n$. The set B consists of $(2/9\delta)^n/2$ cubes of volume $(\delta/2)^n$. Thus, the volume of B is bounded below by some absolute positive constant, and the result follows. **q.e.d.**

Let us now define $\mu(\mathbf{x}) = \tilde{\mu}(\tilde{\mathbf{x}})$ for all $\tilde{\mathbf{x}} \in B$, where $\tilde{\mathbf{x}}$ is chosen so that $\|\mathbf{x} - \tilde{\mathbf{x}}\|_\infty \leq \delta/4$, and we let $\mu(\mathbf{x}) = 0$ for $\mathbf{x} \notin B$.

Lemma 3.5: For every $\mathbf{x} \in B$, we have

$$\int_{[2/3, 1]^n} \hat{P}(\mathbf{y} \mid \mathbf{x}, \mu(\mathbf{x})) \, d\mathbf{y} \geq k\delta,$$

where k is an absolute positive constant.

Proof: For every $\tilde{\mathbf{y}} \in \tilde{S}$, let $Y(\tilde{\mathbf{y}}) = \{\mathbf{y} \in S \mid \|\mathbf{y} - \tilde{\mathbf{y}}\|_\infty \leq \delta/4\}$. Fix some $\mathbf{x} \in B$ and let $\tilde{\mathbf{x}}$ be an element of \tilde{S}_{BAD} such that $\|\mathbf{x} - \tilde{\mathbf{x}}\|_\infty \leq \delta/4$. We have

$$\begin{aligned} \int_{[2/3, 1]^n} \hat{P}(\mathbf{y} \mid \mathbf{x}, \mu(\mathbf{x})) \, d\mathbf{y} &\geq \sum_{\tilde{\mathbf{y}} \in Q^H(\tilde{\mathbf{x}})} \int_{Y(\tilde{\mathbf{y}})} E_{\tilde{\mathbf{y}}, \tilde{\mathbf{x}}, \tilde{\mu}(\tilde{\mathbf{x}})}(\mathbf{y} \mid \mathbf{x}, \mu(\mathbf{x})) \, d\mathbf{y} \\ &\geq \sum_{\tilde{\mathbf{y}} \in Q^H(\tilde{\mathbf{x}})} \int_{Y(\tilde{\mathbf{y}})} \left(\frac{\delta}{2} - \max\{\|\mathbf{x} - \tilde{\mathbf{x}}\|_\infty, \|\mathbf{y} - \tilde{\mathbf{y}}\|_\infty\} \right) \, d\mathbf{y} \\ &\geq \sum_{\tilde{\mathbf{y}} \in Q^H(\tilde{\mathbf{x}})} \frac{\delta}{4} \int_{Y(\tilde{\mathbf{y}})} \, d\mathbf{y}. \end{aligned}$$

The set $Y(\tilde{y})$ is a cube of volume $(\delta/2)^n$, the cardinality of the set $Q^H(\tilde{x})$ is $(2/9\delta)^n/2$, and the result follows. **q.e.d.**

We now estimate the cost $J_\pi(x)$ which is incurred if policy $\pi = (\mu, \mu, \dots)$ is used, for the case where $x \in [0, 1/3]^n$. The corresponding Markov process x_t^π evolves as follows. Whenever $x_t^\pi \in [0, 1/3]^n$, there is at least probability h that the next state belongs to the set B and there is a further probability of at least $k\delta$ that the state after one more transition is in the set $[2/3, 1]^n$. Once the latter set is entered, the state stays forever in that set. We therefore have

$$\Pr(x_t^\pi \in [0, 1/3]^n) \leq (1 - kh\delta)^{t-1}, \quad \forall t \geq 1.$$

Since the cost is 1 on the set $[0, 1/3]^n$ and 0 on the set $[2/3, 1]^n$, we have

$$\begin{aligned} J_\pi(x) &= \sum_{t=0}^{\infty} \alpha^t \Pr(x^\pi(t) \in [0, 1/3]^n) \\ &\leq 1 + \sum_{t=1}^{\infty} \alpha^t (1 - kh\delta)^{t-1} \\ &= 1 + \frac{\alpha}{1 - \alpha(1 - kh\delta)} \\ &= \frac{1 + \alpha kh\delta}{1 - \alpha(1 - kh\delta)}, \quad \forall x \in [0, 1/3]^n. \end{aligned} \tag{3.27}$$

The optimal cost function \hat{J}^* of the perturbed instance satisfies $\hat{J}^* \leq \hat{J}_\pi$ and, using Eq. (3.25), we obtain

$$J^*(x) - \hat{J}^*(x) \geq \frac{1}{1 - \alpha} - \frac{1 + \alpha kh\delta}{1 - \alpha(1 - kh\delta)} = \frac{\alpha^2 kh\delta}{(1 - \alpha)(1 - \alpha(1 - kh\delta))}, \quad \forall x \in [0, 1/3]^n. \tag{3.28}$$

Notice that the class \mathcal{P}_{prob} contains the class \mathcal{P}_{mix} . For this reason, the particular algorithm being considered here is also a correct algorithm for the problem \mathcal{P}_{mix} . In particular, all of the intermediate results in the proof of Theorem 3.1 apply to the algorithm we are considering. We can therefore use Eq. (3.21) and conclude that $\delta \leq d(1 - \alpha)\epsilon$, where d is an absolute constant. (Actually, the definition of δ is somewhat different in the two proofs, but this only affects the absolute constant d .) This implies that for $\epsilon \leq 1/(khd)$, we have $\delta \leq (1 - \alpha)/(kh)$ or $1 - kh\delta \geq \alpha$. Using this inequality in Eq. (3.28), together with the property $\alpha \geq 1/2$, we obtain

$$J^*(x) - \hat{J}^*(x) \geq \frac{\alpha^2 kh\delta}{(1 - \alpha)(1 - \alpha^2)} \geq \frac{(kh\delta)/4}{2(1 - \alpha)^2} = \frac{1}{8} \frac{kh\delta}{(1 - \alpha)^2}.$$

This inequality is similar to inequality (3.20) in the proof of Theorem 3.1, except that $1 - \alpha$ has been replaced by $(1 - \alpha)^2$. The rest of the argument is the same, except for certain constant factors, and that $1 - \alpha$ is replaced throughout by $(1 - \alpha)^2$. **Q.E.D.**

Theorem 3.3: (Lower bound under Assumptions 2.1, 2.3, 2.4) For every m, n , there is a choice of K and ρ such that

$$C_{sub}(\epsilon, \alpha) = \Omega \left(\frac{1}{((1 - \alpha)^2 \epsilon)^{2n+m}} \right).$$

Proof: The proof is almost identical to the proof of Theorem 3.2, and for this reason, we argue informally. For convenience, let the state space S be the set $[0, 1/3]^n$, instead of $[0, 1]^n$, and let $P(y | x, u)$ be defined on that set as in proof of Theorem 3.2. Then, P is a probability measure on the set $[0, 1/3]^n$ and the corresponding function J^* is identically equal to $1/(1 - \alpha)$. Notice that P satisfies Assumption 2.3. Let \hat{P} be as in the proof of Theorem 3.2, except that it is defined only for $x, y \in [0, 1/3]^n$. For this reason, \hat{P} is now a subprobability measure. The function \hat{J}^* for the current problem is equal to the optimal expected discounted cost until the termination of the stochastic process. However, the process considered here terminates exactly when the process considered in the proof of Theorem 3.2 makes a transition from $[0, 1/3]^n$ to the zero-cost set $[2/3, 1]^n$. For this reason, the function \hat{J}^* is the same as the function J^* in the proof of Theorem 3.2, and the result follows with the same reasoning. **Q.E.D.**

Remarks:

1. Suppose that we replace the correctness requirement $\|J - J^*\|_\infty \leq \epsilon$ by the requirement $\|J - J^*\|_p \leq \epsilon$, where $1 \leq p < \infty$ and $\|\cdot\|_p$ is the usual L_p norm. Then, Theorems 3.1–3.3 remain true, with exactly the same proofs. The reason, is that in all of our proofs we have constructed our perturbed instances so that $J^*(x) - \hat{J}^*(x)$ is “large” on a set whose measure is bounded below by an absolute constant [cf. Eq. (3.20) or Eq. (3.28)]. But this implies that $J^* - \hat{J}^*$ is also large when measured by the L_p -norm and the proofs remain valid, except that certain constants have to be changed.
2. The lower bounds of Theorem 3.3 can also be proved for all values of the constants K and ρ . The proof is similar except that we should let $P(y|x, u) = 3^{-n}$ for all $(y, x, u) \in S \times S \times U$, so that P satisfies the Lipschitz continuity assumption for any value of K . Furthermore, the perturbing pyramids should be multiplied by a factor that ensures that their Lipschitz constant is less than K and that Assumption 2.3 is not violated.
3. We are not able to establish the lower bound of Theorem 3.2 for an arbitrary choice of K . There is a simple reason for that: if K is taken very small, then Assumption 2.3 is automatically satisfied and the best provable lower bound is the one in Theorem 3.1.
4. Given some $\epsilon > 0$, we say that a function $\mu : S \mapsto U$ is ϵ -optimal if $J^*(x) \leq J_\pi(x) \leq J^*(x) + \epsilon$, for all $x \in S$, where $\pi = (\mu, \mu, \dots)$. If an ϵ -optimal function μ is available, then the function J^* is automatically determined within an error of ϵ , the error being measured according to the norm $\|\cdot\|_\infty$. Thus, the number of oracle queries needed for computing an ϵ -optimal function μ is at least as large as the number of queries needed to determine J^* within ϵ . It follows that the lower bounds of Theorems 3.1–3.3 also apply (under their respective assumptions) to the

computation of ϵ -optimal functions μ .

4. DISCUSSION

The lower bounds of Section 3 agree with the upper bounds of Section 2. Thus, we have completely characterized the information-based complexity of the approximate computation of J^* . This leaves the further question of evaluating the total complexity of approximating J^* , when arithmetic computations are taken into account. This issue is addressed in [CT]. In particular, [CT] introduces a multigrid version of the iterative algorithm $J := TJ$ and shows that the total number of arithmetic operations and comparisons is

$$O\left(\frac{1}{((1-\alpha)\epsilon)^{2n+m}}\right), \quad (4.1)$$

for the problem \mathcal{P}_{mix} , and

$$O\left(\frac{1}{((1-\alpha)^2\epsilon)^{2n+m}} \cdot \frac{1}{|\log \alpha|}\right) = O\left(\frac{1}{((1-\alpha)^2\epsilon)^{2n+m}} \cdot \frac{1}{1-\alpha}\right), \quad (4.2)$$

for the problems \mathcal{P}_{prob} and \mathcal{P}_{sub} . Thus, for problem \mathcal{P}_{mix} , we have an optimal algorithm. For the problems \mathcal{P}_{prob} and \mathcal{P}_{sub} , we are within a factor of $O(1/(1-\alpha))$ from the optimum. One might wish to close this gap but the prospects are not particularly bright because a) there are no effective methods for proving lower bounds tighter than those provided by the information-based approach and, b) it can be shown [C] that no algorithm in a certain family of multigrid methods can have complexity better than the one provided by Eq. (4.2).

We expect that our results can be extended to the case where bounds are imposed on second derivatives (more generally, derivatives of order r) of the functions P and g . Of course, the bounds should change, with the exponent $2n+m$ being replaced by a lower exponent, depending on r .

As mentioned in the introduction, the case where the functions g and P do not depend on u (equivalently, the case where U is a singleton) makes the equation $J = TJ$ a linear Fredholm equation of the second kind. Our proofs and our results remain true, provided that the exponent m in our bounds is replaced by 0. In particular, if we let $n = 1$, our results agree with the results of [W].² Our results are different from those of [W] in a number of respects:

- a) We are not limited to the one-dimensional case.
- b) We quantify the dependence of the complexity on the parameter α , which is a measure of the

2. In [W] the Lipschitz continuity assumption is replaced by a continuous differentiability assumption. Since the latter is a more restrictive assumption, our lower bounds do not apply to that case. However, the arguments in our proofs can be modified, by “smoothing” the corners and the edges of our “pyramids” and therefore the same lower bounds hold for the continuously differentiable case.

ill-conditioning of the problem.

c) On the other hand, unlike [W], we do not study the dependence of the complexity on the smoothness properties (e.g. bounds on higher derivatives) of the functions P and g .

REFERENCES

- [B1] D.P. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, Englewood Cliffs, New Jersey, 1987.
- [B2] D.P. Bertsekas, "Convergence of Discretization Procedures in Dynamic Programming", *IEEE Transactions on Automatic Control*, Vol. AC-20, 1975, pp. 415-419.
- [BS] D.P. Bertsekas and S.E. Shreve, *Stochastic Optimal Control: The Discrete Time Case*, Academic Press, New York, 1978.
- [C] C.-S. Chow, Ph.D. Thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1989 (in preparation).
- [CT] C.-S. Chow and J.N. Tsitsiklis, "An Optimal Multigrid Algorithm for Discrete-Time Stochastic Control", Technical Report LIDS-P-1864, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, Massachusetts, April 1989. (A preliminary version of this paper was presented at the 27th IEEE Conference on Decision and Control.)
- [HTW] R. Hartley, L. C. Thomas, and D. J. White (Eds.), *Recent Developments in Markov Decision Processes*, Academic Press, New York, 1980.
- [NY] A.S. Nemirovsky and D.B. Yudin, *Problem Complexity and Method Efficiency in Optimization*, J. Wiley & Sons, 1983.
- [PT] C.H. Papadimitriou and J.N. Tsitsiklis, "The Complexity of Markov Decision Processes", *Mathematics of Operations Research*, Vol. 12, 3, 1987, pp. 441-450.
- [TWW] J.F. Traub, G.W. Wasilkowski, and H. Wozniakowski, *Information-Based Complexity*, Academic Press, 1988.
- [W] A. G. Werschulz, "What is the Complexity of the Fredholm Problem of the Second Kind?", *Journal of Integral Equations*, Vol. 9, 1985, pp. 213-241.
- [Wh] W. Whitt, "Approximations of Dynamic Programs I", *Mathematics of Operations Research*, Vol. 3, 1978, pp. 231-243.