# CORC Technical Report TR-2004-10 <br> Ambiguous chance constrained problems and robust optimization* 

E. Erdoğan ${ }^{\dagger} \quad$ G. Iyengar ${ }^{\ddagger}$

September 7, 2004


#### Abstract

In this paper we study ambiguous chance constrained problems where the distributions of the random parameters in the problem are themselves uncertain. We primarily focus on the special case where the uncertainty set $\mathcal{Q}$ of the distributions is of the form $\mathcal{Q}=\left\{\mathbb{Q}: \rho_{p}\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq\right.$ $\beta\}$, where $\rho_{p}$ denotes the Prohorov metric. The ambiguous chance constrained problem is approximated by a robust sampled problem where each constraint is a robust constraint centered at a sample drawn according to the central measure $\mathbb{Q}_{0}$. The main contribution of this paper is to show that the robust sampled problem is a good approximation for the ambiguous chance constrained problem with high probability. This result is established using the Strassen-Dudley Representation Theorem that states that when the distributions of two random variables are close in the Prohorov metric one can construct a coupling of the random variables such that the samples are close with high probability. We also show that the robust sampled problem can be solved efficiently both in theory and in practice.


## 1 Introduction

A large class of decision problems in finance and engineering can be formulated as optimization problems of the form

$$
\begin{align*}
\min & \mathbf{c}^{T} \mathbf{x} \\
\text { s.t. } & f(\mathbf{x}, \mathbf{h})=\max _{1 \leq i \leq l}\left\{f_{i}(\mathbf{x}, \mathbf{h})\right\} \leq 0,  \tag{1}\\
& \mathbf{x} \in \mathcal{X},
\end{align*}
$$

where $\mathbf{x} \in \mathcal{X}$ is the decision vector, $\mathcal{X} \subseteq \mathbf{R}^{n}$ is a closed convex set, $\mathbf{h} \in \mathbf{R}^{m}$ are problem parameters and each of the functions $f_{i}(\mathbf{x}, \mathbf{h}): \mathcal{X} \times \mathcal{H} \mapsto \mathbf{R}$ are convex in $\mathbf{x}$ for a fixed $\mathbf{h}$. We assume, without loss of generality, that the objective is linear and independent of $\mathbf{h}$.

[^0]The deterministic optimization approach to solving optimization problems computes a solution to (1) assuming that the problem parameters $\mathbf{h}$ are known and fixed. In practice, however, the parameters $\mathbf{h}$ are the result of some measurement or estimation process, and are, therefore, never known for certain. This uncertainty is of serious concern in applications because solutions of optimization problems are often very sensitive to fluctuations in the problem parameters. This phenomenon is well documented in several different application areas $[3,26]$.

Recently robust optimization has emerged as an attractive optimization framework for reducing the sensitivity of the optimal solution to perturbations in the parameter values. In this framework, the uncertain parameters $\mathbf{h}$ are assumed to belong to a bounded uncertainty set $\mathcal{H}$ and the robust optimization problem corresponding to the nominal problem (1) is given by

$$
\begin{align*}
\min & \mathbf{c}^{T} \mathbf{x} \\
\text { s.t. } & f(\mathbf{x}, \mathbf{h}) \leq 0, \quad \forall \mathbf{h} \in \mathcal{H}  \tag{2}\\
& \mathbf{x} \in \mathcal{X}
\end{align*}
$$

This framework was introduced in Ben-Tal and Nemirovski [4, 5, 6]. There is also a parallel literature on robust formulations of optimization problems originating from robust control [18, 20]. In many applications the uncertainty set $\mathcal{H}$ is given by the confidence region around the point estimates of the parameter allowing one to provide probabilistic guarantees on the performance of the optimal solution of the robust problem [26]. The robust problem (2) is solved by reformulating the semi-infinite constraints, $f(\mathbf{x}, \mathbf{h}) \leq 0$, for all $\mathbf{h} \in \mathcal{H}$, as a finite collection of constraints. Such a reformulation is only possible when the uncertainty set $\mathcal{H}$ and the function $f(\mathbf{x}, \mathbf{h})$ satisfy some regularity conditions. See $[4,6,7]$ for robust formulations that can be solved efficiently. Even when the reformulation is possible, the resulting problem is typically harder than the nominal problem (1) ([7] proposes a new framework wherein the robust problem remains in the same complexity class as the nominal problem). In general, however, the robust problem is intractable.

Another criticism of the robust framework is that it gives the same "weight" to all perturbations $\mathbf{h} \in \mathcal{H}$. Also, in certain applications one might have the flexibility of violating the constraints corresponding to a small fraction of the set $\mathcal{H}$. An alternative optimization framework that mitigates this criticism to some extent is called chance-constrained optimization. In this framework, one assumes the parameters are $\mathbf{h}$ are distributed according to a known distribution $\mathbb{Q}$ on $\mathcal{H}$, and replaces the nominal problem (1) by the following chance-constrained problem

$$
\begin{align*}
\min & \mathbf{c}^{T} \mathbf{x} \\
\text { s.t. } & \mathbf{x} \in \mathcal{X}_{\epsilon}(\mathbb{Q}) \tag{3}
\end{align*}
$$

where

$$
\begin{equation*}
\mathcal{X}_{\epsilon}(\mathbb{Q})=\{\mathbf{x} \in \mathcal{X}: \mathbb{Q}(\mathbf{H}: f(\mathbf{x}, \mathbf{H})>0) \leq \epsilon\} \tag{4}
\end{equation*}
$$

for some $0<\epsilon<1$. The parameter $\epsilon$ controls the probability that the optimal solution of (3) violates the constraints - as $\epsilon \downarrow 0$ the chance-constrained problem starts to resemble the robust
problem (2). Although chance-constrained problems have a long history dating back to at least the work of Charnes and Cooper [12], they have not found wide applicability. This is primarily because computing the optimal solution for chance-constrained problems is extremely hard. To begin with just evaluating $\mathbb{Q}(\mathbf{H}: f(\mathbf{x}, \mathbf{H})>0)$ involves a multidimensional integral that becomes hard as the number of parameters grows. Moreover, even if the function $f(\mathbf{x}, \mathbf{h})$ is convex (or even linear) in $\mathbf{x}$ the feasible set $\mathcal{X}_{\epsilon}(\mathbb{Q})$ of (3) is not convex. A detailed discussion of the chance-constrained programs and, more generally, stochastic programs can be found in [35].

Recently, Calafiore and Campi $[10,11]$ and de Farias and Van Roy [14] independently proposed tractable approximations to (3) based on constraint sampling and statistical learning techniques. In this approach, one approximates chance-constrained problem (3) by the following sampled problem

$$
\begin{align*}
\min & \mathbf{c}^{T} \mathbf{x} \\
\text { s.t. } & f\left(\mathbf{x}, \mathbf{H}_{i}\right) \leq 0, \quad i=1, \ldots, N  \tag{5}\\
& \mathbf{x} \in \mathcal{X}
\end{align*}
$$

where $\mathbf{H}_{i}, i=1, \ldots, N$, are $N$ independent, identically distributed (IID) samples from the distribution $\mathbb{Q}$. de Farias and Van Roy [14] consider the special case where $f(\mathbf{x}, \mathbf{h})=\mathbf{h}^{T} \mathbf{x}+h_{0}$ and use results from Computational Learning Theory $[1,30,43]$ to show that for all $N \geq \frac{4 n}{\epsilon} \ln \left(\frac{12}{\epsilon}\right)+\frac{4}{\epsilon} \ln \left(\frac{2}{\delta}\right)$, the feasible set of the sampled problem (5) is contained in $\mathcal{X}_{\epsilon}(\mathbb{Q})$ with probability at least $1-\delta$. Thus, in this sampling based method there are two possible sources of errors: with probability $\delta$, the feasible set of (5) (and consequently, the optimal solution of (5)) may not be contained in $\mathcal{X}_{\epsilon}(\mathbb{Q})$; and, in event that this is not the case, the feasible points of (5) can still violate the constraint $f(\mathbf{x}, \mathbf{H}) \leq 0$ with a probability $\epsilon$. The analysis in [14] can be extended to general $f(\mathbf{x}, \mathbf{h})$ (see Section 3 for details). Calafiore and Campi [10,11] consider general convex functions $f(\mathbf{x}, \mathbf{h})$ and show that for $N \geq \frac{2 n}{\epsilon} \ln \left(\frac{12}{\epsilon}\right)+\frac{2}{\epsilon} \ln \left(\frac{2}{\delta}\right)+2 n$, the optimal solution of the sampled problem (5) is feasible for (3) with probability at least $1-\delta$. On the one hand, this bound is weak in the sense that it is only valid for the optimal solution, and not the entire feasible set. On the other hand, the number of samples required to ensure that the optimal solution is feasible for (3) with high probability can be orders of magnitude lower. The result in $[10,11]$ is proved using a fundamental fact that the optimal solution of a convex program is "supported" by at most $n$ constraints. We will briefly review this work in Section 3.3. Recently, Nemirovski and Shapiro [33, 32] established logarithmically separated upper and lower bounds on the number of samples required to approximate a chance constrained problem when the measure $\mathbb{Q}$ has well defined moment generating function.

Although the bounds on the sample size $N$ are distribution-free in the sense that they do not depend on the underlying measure $\mathbb{Q}$, in order to construct the sampled problem (5) one requires samples from this probability measure. Also, there is an implicit assumption that the distribution $\mathbb{Q}$ of the random parameters $\mathbf{H}$ is fixed. A major criticism raised against chance constrained problems and, more generally, stochastic programs is that, in practice, the measure is never known exactly. Just as the point estimates for the parameters, the distribution $\mathbb{Q}$ is also estimated from data or
measurements, and is, therefore, known only to within some error, i.e. the measure $\mathbb{Q} \in \mathcal{Q}$ where $\mathcal{Q}$ is a set of measures. Since our primary interest in the chance constrained problem (3) was to use it as an approximation (or even a refinement) of the robust problem (2), the natural problem to consider when the measure $\mathbb{Q}$ is uncertain is given by

$$
\begin{array}{cl}
\min & \mathbf{c}^{T} \mathbf{x} \\
\text { s.t. } & \mathbf{x} \in \overline{\mathcal{X}}_{\epsilon}, \tag{6}
\end{array}
$$

where

$$
\begin{equation*}
\overline{\mathcal{X}}_{\epsilon}=\{\mathrm{x} \in \mathcal{X}: \mathbb{Q}(\mathbf{H}: f(\mathrm{x}, \mathbf{H})>0) \leq \epsilon, \forall \mathbb{Q} \in \mathcal{Q}\} . \tag{7}
\end{equation*}
$$

We will call (6) an ambiguous chance-constrained problem. A problem of the form (6) has two sources of uncertainty: the distribution $\mathbb{Q}$ of the parameter $\mathbf{h}$ is uncertain, and, given a measure $\mathbb{Q}$, the particular realization of the parameter $\mathbf{h}$ is also uncertain. In the decision theory literature the uncertainty in the distribution is referred to as ambiguity, and hence the name for the problem.

Modeling ambiguity and its consequence has been receiving attention in several different fields. The minimax formulation has a long history in stochastic programing [44, 8, 16, 17, 29, 40, 38, 39]. Ruszczynski and Shapiro [36] show the equivalence between minimax stochastic programming and minimizing a convex risk measure [2, 23] of the second-stage cost. [37] extends the minimax approach to a multiperiod setting. The study of ambiguity in Economics began with the work of Gilboa and Schmeidler [25]. This work was extended to a multiperiod setting by Hansen and Sargent [27] and Epstein and his co-authors [13, 21, 22]. Ambiguity in the context of Markov decision processes was independently investigated by Iyengar [28] and El Ghaoui and Nilim [19].

The main contributions of this paper are as follows.
(a) We consider uncertainty sets $\mathcal{Q}$ of measures that are of the form $\mathcal{Q}=\left\{\mathbb{Q}: \rho\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \beta\right\}$ where $\rho(\cdot, \cdot)$ denotes a suitable metric between probability measures, i.e. the uncertainty sets are "balls" centered around the central measure $\mathbb{Q}_{0}$. We approximate the ambiguous chanceconstrained problem (6) by a robust sampled problem defined as follows

$$
\begin{align*}
\min & \mathbf{c}^{T} \mathbf{x} \\
\text { s.t. } & f(\mathbf{x}, \mathbf{z}) \leq 0, \quad \forall \mathbf{z} \text { s.t. }\left\|\mathbf{z}-\mathbf{H}_{i}^{0}\right\| \leq \beta, \quad i=1, \ldots, N, \tag{8}
\end{align*}
$$

where $\mathbf{H}_{i}^{0}, i=1, \ldots, N$, denote IID samples drawn according to the central measure $\mathbb{Q}_{0}$ and the norm $\|\cdot\|$ on the $\mathcal{H}$ space is related to the probability metric $\rho(\cdot, \cdot)$ (details are given in Section 4). Results in [7] imply that for a large class of constraint functions $f(\mathbf{x}, \mathbf{h})$ and suitably defined norms $\|\cdot\|$ the robust sampled problem (8) is in the same complexity class as the nominal problem (1).
(b) We combine results from Computational Learning Theory with results for coupling of random variables [42] to compute upper bounds on the number of samples $N$ required to ensure that the feasible set of the robust sampled problem (8) is contained in $\overline{\mathcal{X}}_{\epsilon}$ with high probability. This bound depends on the Vapnik-Chervonenkis (VC) dimension of the function $f(\mathbf{x}, \mathbf{h})$.
(c) We use coupling to extend the results of Calafiore and Campi $[10,11]$ to the ambiguous chance constrained problems, i.e. we compute upper bounds on the number of samples required to ensure that the optimal solution of the robust sampled problem (8) is contained in $\overline{\mathcal{X}}_{\epsilon}$ with high probability. The bound in this case depends on the number of "support" constraints, and is independent of the VC dimension of $f(\mathbf{x}, \mathbf{h})$.

The issue of ambiguity of measures was also raised in [11] where the authors considered a finite uncertainty set $\mathcal{Q}$. They proposed a solution strategy where one samples from all of these measures and showed that the samples from different measure "help" each other. In contrast, we consider the case where $\mathcal{Q}$ is uncountably infinite and we draw samples from only the central measure $\mathbb{Q}_{0}$.

The rest of this paper is structured as follows. In Section 3 we briefly review the known results for chance constrained problem. The results in this section are not new - they have been included to set the context for our extensions. Section 4 introduces probability metrics, coupling and the Strassen-Dudley Representation Theorem. Section 5 uses this Representation Theorem to establish bounds for ambiguous chance constrained problems. In Section 6 we identify particular classes of functions $f(\cdot, \cdot)$ and norms $\|\cdot\|$ on the parameter space $\mathcal{H}$ that allow the robust sampled problem (8) to be solved efficiently. Section 7 has some concluding remarks.

## 2 Notation

Sets will be denoted by calligraphic letters, e.g. $\mathcal{A}$. For a finite set $\mathcal{A}$, we will denote the size of $\mathcal{A}$ by $|\mathcal{A}|$. All (deterministic) vectors will be denoted by boldface lowercase letters, e.g. x. Random vectors and samples of random vectors will be denoted by boldface uppercase letters, e.g. $\mathbf{H}$, and measures will be denoted by mathematical boldface letters, e.g. $\mathbb{P}$. We will denote that a random vector $\mathbf{H}$ has distribution $\mathbb{Q}$ by $\mathbf{H} \sim \mathbb{Q}$, a $\sigma$-algebra on a space $\mathcal{H}$ by $\mathcal{F}(\mathcal{H})$, and the set of all probability measures on $\mathcal{H}$ by $\mathcal{M}(\mathcal{H})$. We will denote the $n$-th binomial coefficient $\frac{N!}{(N-n)!n!}$ by $\binom{N}{n}$.

## 3 Chance constrained problems and Learning Theory

In this section our goal is to relate the sampled problem (5) to the chance constrained problem (3). We assume that the distribution $\mathbb{Q}$ of the perturbations $\mathbf{H}$ is known and fixed. Let $\mathbf{H}_{1, N}=$ $\left\{\mathbf{H}_{1}, \mathbf{H}_{2}, \ldots, \mathbf{H}_{N}\right\}$ denote $N$ IID samples of the random vector $\mathbf{H} \sim \mathbb{Q}$. Then the feasible set of the sampled problem (5) is given by

$$
\begin{equation*}
\mathcal{Y}\left[\mathbf{H}_{1, N}\right]=\left\{\mathbf{x} \in \mathcal{X}: f\left(\mathbf{x}, \mathbf{H}_{i}\right) \leq 0, i=1, \ldots, N\right\} \tag{9}
\end{equation*}
$$

In the sequel, we will often abbreviate $\mathcal{Y}\left[\mathbf{H}_{1, N}\right]$ as $\mathcal{Y}[N]$ with the understanding that the set $\mathcal{Y}[N]$ is defined using a particular sequence of IID samples $\mathbf{H}_{1, N}$ of length $N$. In the first half of this section we appropriately interpret concepts from Computational Learning Theory [1, 43, 30] to
establish bounds on the number of samples $N$ required for $\mathcal{Y}[N]$ to be a good approximation for the feasible set $\mathcal{X}_{\epsilon}=\{\mathbf{x}: \mathbb{Q}(\mathbf{H}: f(\mathbf{x}, \mathbf{H}) \leq 0) \leq \epsilon\}$ of the chance constrained problem (3). Next, we briefly summarize the result in $[10,11]$ that provides a bound for the number of samples required to ensure that the optimal solution of the sampled problem (5) is contained in $\mathcal{X}_{\epsilon}$ with high probability. As mentioned in Section 1, the results in this section are not new - they simply provide the context for the new results in Section 5.

### 3.1 Decision vectors, concepts and Vapnik-Chervonenkis (VC) dimension

With each decision vector $\mathbf{x} \in \mathcal{X}$ we associate the concept or classification

$$
\mathcal{C}_{x}=\{\mathbf{h} \in \mathcal{H}: f(\mathbf{x}, \mathbf{h}) \leq 0\} .
$$

Let $\mathcal{C}_{f}=\left\{\mathcal{C}_{x}: \mathbf{x} \in \mathcal{X}\right\}$ denote the class of all concepts induced on $\mathcal{H}$ as the decision vector $\mathbf{x}$ runs over the set $\mathcal{X}$. Thus, the set $\mathcal{X}_{\epsilon}=\left\{\mathbf{x} \in \mathcal{X}: \mathbb{Q}\left(\mathbf{H} \notin \mathcal{C}_{x}\right) \leq \epsilon\right\}$. To apply the results from Learning Theory to the problem at hand we pretend that our goal is to learn the set $\mathcal{X}_{0}$ of concepts $\mathcal{C}_{x}$ that cover $\mathcal{H}$ with probability 1 , i.e. $\mathcal{X}_{0}=\left\{\mathbf{x}: \mathbb{Q}\left(\mathbf{H} \in \mathcal{C}_{x}\right)=1\right\}$. Since Learning algorithms only have access to a finite number of samples of the random variable $\mathbf{H}$, it is impossible to learn the concepts in $\mathcal{X}_{0}$; instead any such algorithm will have to be satisfied with learning a concept with a small error $\epsilon$, i.e. a concept $\mathcal{C}_{x}$ with $\mathbb{Q}\left(\mathbf{H} \notin \mathcal{C}_{x}\right) \leq \epsilon$ or equivalently $\mathbf{x} \in \mathcal{X}_{\epsilon}$. For the particular case considered in this paper, learning $\mathcal{X}_{\epsilon}$ is equivalent to producing a good approximation for the function $f(\mathbf{x}, \cdot)$ using a finite number of samples. Thus, one should expect that the complexity of learning $\mathcal{X}_{\epsilon}$ when the function $f(\mathbf{x}, \mathbf{h})=\mathbf{h}^{T} \mathbf{x}$ is linear, or equivalently the associated concept $\mathcal{C}_{x}$ is a half space, should be smaller than the complexity of learning $\mathcal{X}_{\epsilon}$ when the function $f(\cdot, \cdot)$ is nonlinear. Learning Theory quantifies the complexity of a concept class $\mathcal{C}_{f}$ by its Vapnik-Chervonenkis (VC) dimension [43].

Let $\mathcal{S}=\left\{\mathbf{s}_{1}, \ldots, \mathbf{s}_{N}\right\} \subset \mathcal{H}$ denote a finite subset of $\mathcal{H}$ with $|S|=N$. Define

$$
\begin{equation*}
\Pi_{f}(\mathcal{S})=\left\{\left(\mathbf{1}_{\mathcal{C}_{x}}\left(\mathbf{s}_{1}\right), \ldots, \mathbf{1}_{\mathcal{C}_{x}}\left(\mathbf{s}_{N}\right)\right): \mathbf{x} \in \mathcal{X}\right\} \tag{10}
\end{equation*}
$$

where $\mathbf{1}_{\mathcal{C}}(\cdot)$ denotes the characteristic function of the set $\mathcal{C}$. The set $\Pi_{f}$ is the set of dichotomies or behaviors induced by the concept class $\mathcal{C}_{f}$, or equivalently the function $f(\cdot, \cdot)$. From (10), it is clear that the number of elements $\left|\Pi_{f}(\mathcal{S})\right| \leq 2^{N}$. We say that a set $\mathcal{S}$ is shattered by the concept class $\mathcal{C}_{f}$ if $\Pi_{f}(\mathcal{S})=\{0,1\}^{N}$, or equivalently $\left|\Pi_{f}(\mathcal{S})\right|=2^{N}$. Note that if a set $\mathcal{S}$ is shattered by the concept class $\mathcal{C}_{f}$ it does not yield any information about the concept class. Thus, the size of largest shattered set is a measure of the complexity of the concept class $\mathcal{C}_{f}$.

Definition 1 (VC dimension of $f(\cdot, \cdot)$ ) The $V C$ dimension $d_{f}$ of the function $f(\cdot, \cdot)$ is the cardinality of the largest set $\mathcal{S} \subset \mathcal{H}$ that is shattered by the concept class $\mathcal{C}_{f}$, i.e

$$
\begin{align*}
d_{f} & =\sup \left\{|\mathcal{S}|: \Pi_{f}(\mathcal{S})=\{0,1\}^{N}\right\}, \\
& =\sup \left\{|\mathcal{S}|:\left|\Pi_{f}(\mathcal{S})\right|=2^{N}\right\} . \tag{11}
\end{align*}
$$

In the sequel we will find it convenient to work with the growth function $\pi_{f}(N)$ defined as follows.

$$
\begin{equation*}
\pi_{f}(N)=\max \left\{\left|\Pi_{f}(\mathcal{S})\right|:|\mathcal{S}|=N\right\} \tag{12}
\end{equation*}
$$

The growth function $\pi_{f}$ is another measure of the complexity of the concept class: the faster this function grows, the more behaviors on sets of size $m$ that can be realized by $\mathcal{C}_{f}$; consequently, the less is the information that this finite set conveys about the class $\mathcal{C}_{f}$. A surprising and fundamental result in Computational Learning Theory asserts that if the VC dimension $d_{f}<\infty$, the growth function $\pi_{f}(N)$ is bounded by a polynomial in $N$ of degree $d_{f}$.

Proposition 1 (Sauer's Lemma [9, 1, 30]) Suppose the VC dimension $d_{f}$ of the function $f(\cdot, \cdot)$ is finite. Then

$$
\begin{equation*}
\pi_{f}(N) \leq 1+\binom{N}{1}+\binom{N}{2}+\ldots+\binom{N}{d_{f}} \leq\left(\frac{e N}{d_{f}}\right)^{d_{f}} \tag{13}
\end{equation*}
$$

where e denotes the base of natural logarithm.
In this paper we assume that the VC dimension $d_{f}<\infty$. This is not a very restrictive assumption since many functions $f(\cdot, \cdot)$ used in practice have finite VC dimension.

Proposition 2 Let $d_{f}$ denote the $V C$ dimension of the function $f(\cdot, \cdot)$.
(a) $\mathcal{X}=\mathbf{R}^{n}, \mathcal{H}=\left\{\left(h_{0}, \mathbf{h}\right): h_{0} \in \mathbf{R}, \mathbf{h} \in \mathbf{R}^{n}\right\}=\mathbf{R}^{n+1}$ and $f(\mathbf{x}, \mathbf{h})=h_{0}+\mathbf{h}^{T} \mathbf{x}$. Then $d_{f} \leq n$.
(b) $\mathcal{X}=\mathbf{R}^{n}, \mathcal{H}=\left\{(\mathbf{A}, \mathbf{b}, \mathbf{u}, v): \mathbf{A} \in \mathbf{R}^{p \times n}, \mathbf{b}, \mathbf{c} \in \mathbf{R}^{n}, v \in \mathbf{R}\right\}$, and $f(\mathbf{x}, \mathbf{h})=\|\mathbf{A x}+\mathbf{b}\|-\mathbf{u}^{T} \mathbf{x}-v$. Then $d_{f} \leq O\left(n^{2}\right)$.
(c) Suppose the $V C$ dimension of the function $f_{i}(\cdot, \cdot)$ is $d_{i}, i=1, \ldots, l$. Then the $V C$ dimension $d_{f}$ of the function $f(\mathbf{x}, \mathbf{h})=\max _{1 \leq i \leq l}\left\{f_{i}(\mathbf{x}, \mathbf{h})\right\}$ is bounded above by $d_{f} \leq \mathcal{O}\left(10^{l} \max _{1 \leq i \leq l}\left\{d_{i}\right\}\right)$.

Proof: Part (a) is proved on p. 77 in [1] (see also [14]), part (b) is established in [9] and part (c) can be established using techniques in [31].
Part (c) states that the best known bound on the VC dimension of a pointwise maximum of $l$ functions grows exponentially in $l$. Thus, the VC dimension of the concept class induced by constraint function $f(\cdot, \cdot)$ of the nominal problem (1) can be quite large. We will remark on this at the end of the next section.

### 3.2 Learning the chance constrained set $\mathcal{X}_{\epsilon}$

For $\mathbf{x} \in \mathcal{X}$ let $\operatorname{err}(\mathbf{x})=\mathbb{Q}\left(\mathbf{H} \notin \mathcal{C}_{x}\right)$. Thus, $\mathcal{X}_{\epsilon}=\{\mathbf{x} \in \mathcal{X}: \operatorname{err}(\mathbf{x}) \leq \epsilon\}$. The feasible set $\mathcal{Y}[N]$ of the sampled problem (5) is the set of all decision vectors $\mathbf{x}$, or equivalently concepts $\mathcal{C}_{x}$, that are consistent with the given sample $\mathbf{H}_{1, N}$. Intuitively speaking, if the sample size is large enough one would expect that $\mathcal{Y}[N]$ is a good estimate of the set $\mathcal{X}_{\epsilon}$. The next two results make this rigorous.

Lemma 1 Fix $\epsilon>0$. Suppose $\overline{\mathbf{x}} \in \mathcal{X}$ with $\operatorname{err}(\overline{\mathbf{x}})>\epsilon$. Then, for all $N \geq 1$,

$$
\begin{equation*}
\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \overline{\mathbf{x}} \in \mathcal{Y}[N]\right) \leq e^{-\epsilon N} \tag{14}
\end{equation*}
$$

where $\mathbb{Q}^{N}$ denotes the product measure $\mathbb{Q} \times \mathbb{Q} \times \ldots \times \mathbb{Q}$ with $N$ terms.
Proof: Recall that $\mathbf{H}_{1, N}$ are IID samples of the random vector $\mathbf{H} \sim \mathbb{Q}$. Therefore,

$$
\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \overline{\mathbf{x}} \in \mathcal{Y}[N]\right)=(\mathbb{Q}(\mathbf{H}: f(\overline{\mathbf{x}}, \mathbf{H}) \leq 0))^{N} \leq(1-\epsilon)^{N} \leq e^{-\epsilon N}
$$

where the last inequality follows from the fact that $1-\epsilon \leq e^{-\epsilon}$.
Lemma 1 establishes that the probability that a given concept $\mathcal{C}_{x}$ with $\operatorname{err}(\mathbf{x})>\epsilon$ is contained in $\mathcal{Y}[N]$ decays exponentially with the number of samples $N$. Suppose the set $\mathcal{X}$ is finite. Then the union bound implies that $\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \mathcal{Y}[N] \nsubseteq \mathcal{X}_{\epsilon}\right) \leq|X| e^{-\epsilon N} \leq \delta$, for all $N \geq \frac{1}{\epsilon} \log \left(\frac{|X|}{\delta}\right)$, i.e $\mathcal{O}\left(\frac{1}{\epsilon} \log \left(\frac{|X|}{\delta}\right)\right)$ samples are needed to learn $\mathcal{X}_{\epsilon}$ with a probability of error bounded by $\delta$. Since the complexity of learning a concept is determined by the VC dimension of the concept class, we expect that a similar bound should hold with $|X|$ replaced by $\pi_{f}(N)$.

Lemma 2 (Proposition 8.2.3 in [1]) Let $\pi_{f}$ denote the growth function associated with concept class $\mathcal{C}_{f}$ induced by $f(\cdot, \cdot)$. Then, for all $N \geq 8 / \epsilon$,

$$
\begin{equation*}
\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \mathcal{Y}[N] \nsubseteq \mathcal{X}_{\epsilon}\right) \leq 2 \pi_{f}(2 N) 2^{-\epsilon N / 2} . \tag{15}
\end{equation*}
$$

This result and the upper bound (13) imply the following corollary.
Corollary 1 Fix $\epsilon, \delta>0$. Suppose the $V C$ dimension $d_{f}$ of $f(\cdot, \cdot)$ is finite. Then

$$
\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \mathcal{Y}[N] \nsubseteq \mathcal{X}_{\epsilon}\right) \leq \delta,
$$

for all

$$
N \geq \max \left\{\frac{8}{\epsilon},\left(\frac{4 d_{f}}{\epsilon} \log \left(\frac{12}{\epsilon}\right)+\frac{4}{\epsilon} \log \left(\frac{2}{\delta}\right)\right)\right\}
$$

We conclude this section with the following lower bound.
Lemma 3 (Theorem 3.5 in [30]) Suppose the VC dimension $d_{f}$ of the function by $f(\cdot, \cdot)$ is $f$ nite. Then the worst case number of samples required to learn $\mathcal{X}_{\epsilon}$ is $\Omega\left(d_{f} / \epsilon\right)$.

Corollary 1 and Lemma 3 establish that the number of samples $N=\Theta\left(d_{f} / \epsilon\right)$. From Proposition 2 (c) we have that the VC dimension of the constraint $f(\cdot, \cdot)$ in the nominal problem (1) could be as large as $10^{l} \max _{1 \leq i \leq l}\left\{d_{i}\right\}$ where $d_{i}$ is VC dimension of the functions $f_{i}, i=1, \ldots, l$. Thus, the number of samples required to learn $\mathcal{X}_{\epsilon}$ could be prohibitive even for well behaved constraint functions.

### 3.3 Quality of the optimal solution of the sampled problem

In this section the goal is more modest - we want to compute the number of samples required to ensure that only the optimal solution of the sampled problem (5), as opposed to the entire set $\mathcal{Y}[N]$, is feasible for the chance constrained problem (3) with high probability. Calafiore and Campi $[10,11]$ recently showed that $N=\mathcal{O}(n / \epsilon)$ is enough to achieve this goal. In this section we briefly review the results in $[10,11]$.

Let $(P)$ denote the following convex program

$$
\begin{aligned}
\min & \mathbf{c}^{T} \mathbf{x} \\
\text { s.t. } & f_{i}(\mathbf{x}) \leq 0, \quad i=1, \ldots, N,
\end{aligned}
$$

where $f_{i}(\mathbf{x})$ is a convex function of $\mathbf{x}$ for all $i=1, \ldots, N$. Let $\widehat{\mathbf{x}}$ denote the unique optimal solution of $(P)$. Let $\left(P_{k}\right)$ denote the convex program obtained by dropping the $k$-th constraint, $k=1, \ldots, N$, and let $\widehat{\mathbf{x}}_{k}$ denote the unique optimal solution of the problem $\left(P_{k}\right)$. See $[10,11]$ for the case where the optimal solutions are not unique.

Definition 2 (Support constraint) The $k$-th constraint $f_{k}(\mathbf{x}) \leq 0$ is called a support constraint for the problem $(P)$ if $\mathbf{c}^{T} \widehat{\mathbf{x}}_{k}<\mathbf{c}^{T} \widehat{\mathbf{x}}$.

Theorem 3 (Theorem 2 in [11]) The convex program has at most $n$ support constraints.
Lemma 4 Fix $\epsilon>0$. Let $\widehat{\mathbf{x}}$ denote the optimal solution of the sampled problem (5). Then $\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \widehat{\mathbf{x}} \notin \mathcal{X}_{\epsilon}\right) \leq\binom{ N}{n} e^{-\epsilon(N-n)}$.

Proof: The sampled problem (5) is a convex program with $N$ constraints. Let $\mathcal{I} \subseteq\{1, \ldots, N\}$ with $|I|=n$. Let $\mathcal{H}_{\mathcal{I}}^{N}=\left\{\left(\mathbf{h}_{1}, \ldots, \mathbf{h}_{N}\right):\left(\mathbf{h}_{i}\right)_{i \in \mathcal{I}}\right.$ are the support constraints $\}$. Then Theorem 3 implies $\mathcal{H}^{N}=\cup_{\{\mathcal{I} \subseteq\{1, \ldots, N\}:|\mathcal{I}|=n\}} \mathcal{H}_{\mathcal{I}}^{N}$. Thus,

$$
\begin{aligned}
\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \widehat{\mathbf{x}} \notin \mathcal{X}_{\epsilon}\right) & =\sum_{\{\mathcal{I} \subseteq\{1, \ldots, N\}:|\mathcal{I}|=n\}} \mathbb{Q}^{N}\left(\mathbf{H}_{1, N} \in \mathcal{H}_{\mathcal{I}}^{N}: \widehat{\mathbf{x}}_{\mathcal{I}} \notin \mathcal{X}_{\epsilon}\right) \\
& =\sum_{\{\mathcal{I} \subseteq\{1, \ldots, N\}:|\mathcal{I}|=n\}}\left(\mathbb{Q}^{n}\left(\mathbf{H}_{i \in I}: \widehat{\mathbf{x}}_{\mathcal{I}} \notin \mathcal{X}_{\epsilon}\right) \prod_{i \notin \mathcal{I}} \mathbb{Q}\left(\mathbf{H}_{i}: f\left(\widehat{\mathbf{x}}_{\mathcal{I}}, \mathbf{H}_{i}\right) \leq 0 \mid \mathcal{A}_{\mathcal{I}}\right)\right),
\end{aligned}
$$

where $\widehat{\mathbf{x}}_{\mathcal{I}}$ denotes the optimal solution of the sampled problem (5) with only the samples $i \in \mathcal{I}$ present, $\mathcal{A}_{\mathcal{I}}$ is the event $\mathcal{A}_{\mathcal{I}}=\left\{\mathbf{H}_{i \in I}: \widehat{\mathbf{x}}_{\mathcal{I}} \notin \mathcal{X}_{\epsilon}\right\}$ and each probability in the sum can be written as a product because $\mathbf{H}_{1, N}$ are IID samples. Since $\widehat{\mathbf{x}} \notin \mathcal{X}_{\epsilon}$, it follows that $\mathbb{Q}\left(\mathbf{H}_{i}: f\left(\widehat{\mathbf{x}}_{\mathcal{I}}, \mathbf{H}_{i}\right) \leq 0 \mid \mathcal{A}_{\mathcal{I}}\right) \leq$ $(1-\epsilon)$, for all $i \notin \mathcal{I}$. Thus,

$$
\begin{aligned}
\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \widehat{\mathbf{x}} \notin \mathcal{X}_{\epsilon}\right) & \leq(1-\epsilon)^{(N-n)} \sum_{\{\mathcal{I} \subseteq\{1, \ldots, N\}\}:|\mathcal{I}|=n\}} \mathbb{Q}^{n}\left(\mathbf{H}_{i \in I}: \widehat{\mathbf{x}}_{\mathcal{I}} \notin \mathcal{X}_{\epsilon}\right) \\
& \leq\binom{ N}{n}(1-\epsilon)^{(N-n)} \leq\binom{ N}{n} e^{-\epsilon(N-n)} .
\end{aligned}
$$

Lemma 4 immediately implies the following.

Corollary 2 Fix $\epsilon, \delta>0$. Let $\widehat{\mathbf{x}}$ denote the optimal solution of the sampled problem (5). Then

$$
\mathbb{Q}^{N}\left(\mathbf{H}_{1, N}: \widehat{\mathbf{x}} \nsubseteq \mathcal{X}_{\epsilon}\right) \leq \delta,
$$

for all

$$
N \geq \frac{2 n}{\epsilon} \log \left(\frac{1}{\epsilon}\right)+\frac{2}{\epsilon} \log \left(\frac{1}{\delta}\right)+2 n
$$

## 4 Probability metrics and Coupling

In Section 1 we had introduced the following robust chance constrained set (see (7))

$$
\overline{\mathcal{X}}_{\epsilon}=\left\{\mathbf{x} \in \mathcal{X}: \sup _{\mathbb{Q} \in \mathcal{Q}} \mathbb{Q}(\mathbf{H}: f(\mathbf{x}, \mathbf{H})>0) \leq \epsilon\right\},
$$

where $\mathcal{Q}=\left\{\mathbb{Q}: \rho\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \beta\right\}$ for an appropriately chosen metric $\rho$ on the space $\mathcal{M}(\mathcal{H})$ of probability measures on $\mathcal{H}$. Recall that $\mathcal{X} \subseteq \mathbf{R}^{n}, \mathcal{H} \subseteq \mathbf{R}^{m}$ and we denote the norm in $\mathcal{H}$ space by $\|\cdot\|$. In this section we first review properties of some basic probability metrics. Next, we introduce the concept of coupling of random variables that plays an important role in constructing approximations of the robust chance constrained set $\overline{\mathcal{X}}_{\epsilon}$ via samples. In this paper we will be primarily using the Prohorov metric $\rho_{p}$ defined as follows.

$$
\begin{equation*}
\rho_{p}\left(\mathbb{Q}^{1}, \mathbb{Q}^{2}\right)=\inf \left\{\epsilon: \mathbb{Q}^{1}(\mathcal{B}) \leq \mathbb{Q}^{2}\left(\mathcal{B}^{\epsilon}\right)+\epsilon, \forall \mathcal{B} \in \mathcal{F}(\mathcal{H})\right\}, \tag{16}
\end{equation*}
$$

where

$$
\mathcal{B}^{\epsilon}=\left\{\mathbf{x} \in \mathcal{X}: \inf _{\mathbf{z} \in \mathcal{B}}\|\mathbf{x}-\mathbf{z}\| \leq \epsilon\right\} .
$$

Although the definition appears asymmetric, $\rho_{p}$ is a metric. It plays an important role in probability because it metrizes weak convergence. Moreover, $\rho_{p}\left(\mathbb{Q}_{1}, \mathbb{Q}_{2}\right)$ is the minimum distance "in probability" between random variables distributed according to $\mathbb{Q}_{i}, i=1,2$. Some other metrics of interest are as follows.
(a) Wassestein or Kantorovich metric $\rho_{w}$ :

$$
\rho_{w}\left(\mathbb{Q}_{1}, \mathbb{Q}_{2}\right)=\sup \left\{\left|\int_{\mathcal{H}} g(\mathbf{h})\left(\mathbb{Q}_{1}(d \mathbf{h})-\mathbb{Q}_{2}(d \mathbf{h})\right)\right|: g \in C_{1,1}(\mathcal{H})\right\},
$$

where $C_{1,1}(\mathcal{H})$ denotes the set of Lipschitz continuous functions with Lipschitz constant at most 1 .
(b) Total variation metric $\rho_{t v}$ :

$$
\rho_{t v}\left(\mathbb{Q}_{1}, \mathbb{Q}_{2}\right)=\sup \left\{\left|\mathbb{Q}_{1}(\mathcal{B})-\mathbb{Q}_{2}(\mathcal{B})\right|: \mathcal{B} \in \mathcal{F}(\mathcal{H})\right\} .
$$

(c) Hellinger metric $\rho_{h}$ : Let $f_{i}, i=1,2$ denote the densities of measures $\mathbb{Q}_{i}, i=1,2$, with respect to a common dominating measure (e.g. $\left.\mathbb{Q}=\left(\mathbb{Q}_{1}+\mathbb{Q}_{2}\right) / 2\right)$. Then

$$
\rho_{h}\left(\mathbb{Q}_{1}, \mathbb{Q}_{2}\right)=\left(\int_{\mathcal{H}}\left(\sqrt{f_{1}}-\sqrt{f_{2}}\right)^{2} \mathbb{Q}(d \mathbf{h})\right)^{\frac{1}{2}} .
$$

(d) Relative entropy distance $\rho_{e}$ : Let $f_{i}, i=1,2$ denote the densities of measures $\mathbb{Q}_{i}, i=1,2$, with respect to a common dominating measure (e.g. $\left.\mathbb{Q}=\left(\mathbb{Q}_{1}+\mathbb{Q}_{2}\right) / 2\right)$. Then

$$
\rho_{e}\left(\mathbb{Q}_{1}, \mathbb{Q}_{2}\right)=\int_{\mathcal{H}} f_{1}(\mathbf{h}) \log \left(\frac{f_{1}(\mathbf{h})}{f_{2}(\mathbf{h})}\right) d \mathbf{h}
$$

The relative entropy distance $\rho_{e}$ is not a metric because it is not symmetric and does not satisfy the triangle inequality.

The following lemma relates the Prohorov metric $\rho_{p}$ to the other distance functions.
Lemma 5 ([24]) The distances $\rho_{w}, \rho_{h}, \rho_{t v}$ and $\rho_{e}$ are related to the Prohorov metric as follows.
(a) Prohorov and Wasserstein metrics: $\rho_{p}^{2} \leq \rho_{w} \leq(\operatorname{diam}(\mathcal{H})+1) \rho_{p}$, where $\operatorname{diam}(\mathcal{H})=\sup \left\{\| \mathbf{h}_{1}-\right.$ $\left.\mathbf{h}_{2} \|: \mathbf{h}_{i} \in \mathcal{H}, i=1,2\right\}$.
(b) Prohorov and Total variation metrics: $\rho_{p} \leq \rho_{t v}$
(c) Prohorov and Hellinger metrics: $\rho_{p} \leq \rho_{h}$
(d) Prohorov metric and the relative entropy distance: $\rho_{p} \leq \sqrt{\rho_{e} / 2}$

These bounds imply that for any uncertainty set of the form $\mathcal{Q}=\left\{\mathbb{Q}: \rho\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \delta\right\}$, where the metric $\rho$ is given by $\rho_{w}, \rho_{t v}, \rho_{h}$ or $\rho_{e}$, one can choose $\beta(\delta)>0$ such that $\mathcal{Q} \subseteq \tilde{\mathcal{Q}}=\left\{\mathbb{Q}: \rho_{p}\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq\right.$ $\beta(\delta)\}$, i.e. $\tilde{\mathcal{Q}}$ is a conservative approximation of $\mathcal{Q}$. Next we introduce the concept of coupling of random variables and relate it to the probability metrics.

Definition 3 (Coupling of random variables) A random variable $\tilde{\mathbf{X}}$ is said to be a copy or a representation of the random variable $\mathbf{X}$ if and only if they have the same distribution, i.e. $\tilde{\mathbf{X}} \stackrel{D}{=} \mathbf{X}$. A collection of random variables $\left\{\tilde{\mathbf{X}}^{\alpha}: \alpha \in \mathcal{A}\right\}$ defined on a common probability space $(\Omega, \mathcal{F}(\Omega), \mathbb{P})$ is said to be a coupling of the collection $\left\{\mathbf{X}^{\alpha}: \alpha \in \mathcal{A}\right\}$ if and only if $\tilde{\mathbf{X}}^{\alpha} \stackrel{D}{=} \mathbf{X}$, for all $\alpha \in \mathcal{A}$.

Note that only the individual $\tilde{\mathbf{X}}^{\alpha}$ are copies of the individual $\mathbf{X}^{\alpha}$, the whole collection is $\left\{\tilde{\mathbf{X}}^{\alpha}: \alpha \in\right.$ $\mathcal{A}\}$ is not a copy of $\left\{\mathbf{X}^{\alpha}: \alpha \in \mathcal{A}\right\}$, i.e. the joint distribution of $\left\{\tilde{\mathbf{X}}^{\alpha}: \alpha \in \mathcal{A}\right\}$ need not be the same as that of $\left\{\mathbf{X}^{\alpha}: \alpha \in \mathcal{A}\right\}$.

Theorem 4 (Strassen-Dudley) Let $\mathbf{X}^{1} \sim \mathbb{Q}_{1}$ and $\mathbf{X}^{2} \sim \mathbb{Q}_{2}$ be two random variables taking values in $\mathcal{H}$. Suppose $\rho_{p}\left(\mathbb{Q}_{1}, \mathbb{Q}_{2}\right) \leq \beta$. Then there exists a coupling $\left(\tilde{\mathbf{X}}_{1}, \tilde{\mathbf{X}}_{2}\right)$ of $\left(\mathbf{X}_{1}, \mathbf{X}_{2}\right)$ such that

$$
\begin{equation*}
\mathbb{P}\left(\left\|\tilde{\mathbf{X}}_{1}-\tilde{\mathbf{X}}_{2}\right\|>\beta\right) \leq \beta \tag{17}
\end{equation*}
$$

Proof: This result was established by Strassen [41] for complete separable metric spaces and extended to arbitrary separable metric spaces by Dudley [15]. See also Rachev [34].
This result establishes that if two probability measures $\mathbb{Q}_{i}, i=1,2$, are "close" in the Prohorov metric then there exists a coupling $\left(\tilde{\mathbf{X}}_{1}, \tilde{\mathbf{X}}_{2}\right)$ such that the samples are "close" with high probability. This result can be improved if the random variables $\mathbf{X}_{i}, i=1,2$, are bounded w.p.1.

Theorem 5 Let $\mathbf{X}^{1} \sim \mathbb{Q}_{1}$ and $\mathbf{X}^{2} \sim \mathbb{Q}_{2}$ are two random variables taking values in $\mathcal{H}$. Suppose $\rho_{p}\left(\mathbb{Q}_{1}, \mathbb{Q}_{2}\right) \leq \beta$ and $\left\|\mathbf{X}_{i}\right\| \leq R$ a.s., $i=1,2$. Then there exists a coupling $\left(\tilde{\mathbf{X}}_{1}, \tilde{\mathbf{X}}_{2}\right)$ of $\left(\mathbf{X}_{1}, \mathbf{X}_{2}\right)$ such that

$$
\begin{equation*}
\mathbb{E}\left(\left\|\tilde{\mathbf{X}}_{1}-\tilde{\mathbf{X}}_{2}\right\|\right) \leq(1+2 R) \beta \tag{18}
\end{equation*}
$$

where the expectation is with respect to the common probability measure $\mathbb{P}$.
Proof: The Wasserstein metric $\rho_{w}\left(\mathbb{Q}^{1}, \mathbb{Q}^{2}\right)$ between probability measures $\mathbb{Q}^{1}$ and $\mathbb{Q}^{2}$ can be equivalently characterized as follows.

$$
\rho_{w}\left(\mathbb{Q}_{1}, \mathbb{Q}_{2}\right)=\inf \left\{\mathbb{E}\left[\left\|\tilde{\mathbf{X}}^{1}-\tilde{\mathbf{X}}^{2}\right\|\right]: \mathbf{X}^{i} \sim \mathbb{Q}_{i}, i=1,2,\left(\tilde{\mathbf{X}}^{1}, \tilde{\mathbf{X}}^{2}\right) \text { is a coupling of }\left(\mathbf{X}^{1}, \mathbf{X}^{2}\right)\right\} .
$$

Since $\left\|\mathbf{X}_{i}\right\| \leq R$ a.s., one can without loss of generality assume that $\operatorname{diam}(\mathcal{H}) \leq 2 R$. Thus, the bound $\rho_{w} \leq(\operatorname{diam}(\mathcal{H})+1) \rho_{p}$ together with the characterization above, yields the result.

## 5 Robust chance constrained sets

In this section we show how to construct sampling based approximation for the robust chance constrained set

$$
\overline{\mathcal{X}}_{\epsilon}=\left\{\mathbf{x} \in \mathcal{X}: \sup _{\mathbb{Q} \in \mathcal{Q}} \mathbb{Q}(\mathbf{H}: f(\mathbf{x}, \mathbf{H})>0) \leq \epsilon\right\},
$$

where $\mathcal{Q}=\left\{\mathbb{Q}: \rho_{p}\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \beta\right\}$, and $\rho_{p}$ denotes the Prohorov metric. Note that the bounds in Lemma 5 imply that one can conservatively approximate an uncertainty set defined in terms of any of the metrics discussed in Section 4 by a set defined in terms of the Prohorov metric. The main results of this section are the robust analogs of Lemma 1, Lemma 2 and Lemma 4.

In this section we define err(x) as follows.

$$
\begin{equation*}
\operatorname{err}(\mathbf{x})=\sup _{\mathbb{Q} \in \mathcal{Q}} \mathbb{Q}(\mathbf{H}: f(\mathbf{x}, \mathbf{H})>0) \tag{19}
\end{equation*}
$$

Thus, $\overline{\mathcal{X}}_{\epsilon}=\{\mathbf{x} \in \mathcal{X}: \operatorname{err}(\mathbf{x}) \leq \epsilon\}$. Let $\mathbf{H}_{1, N}^{0}=\left\{\mathbf{H}_{1}^{0}, \ldots, \mathbf{H}_{N}^{0}\right\}$ denote $N$ IID samples drawn according to the central probability measure $\mathbb{Q}_{0}$. Let $\mathcal{Y}[N, \beta]$ denote the set

$$
\begin{equation*}
\mathcal{Y}[N, \beta]=\left\{\mathbf{x}: f(\mathbf{x}, \mathbf{z}) \leq 0, \forall \mathbf{z} \text { s.t. }\left\|\mathbf{z}-\mathbf{H}_{i}^{0}\right\| \leq \beta, i=1, \ldots, N\right\} \tag{20}
\end{equation*}
$$

As in Section 3, the set $\mathcal{Y}[N, \beta]$ is defined using a particular sequence of IID samples $\mathbf{H}_{1, N}^{0}$ of length $N$ drawn according to the measure $\mathbb{Q}_{0}$.

Lemma 6 Fix $\epsilon>0$. Suppose $\overline{\mathbf{x}} \in \mathcal{X}$ with $\operatorname{err}(\overline{\mathbf{x}})>\epsilon$. Then, for all $N \geq 1$,

$$
\begin{equation*}
\mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N}^{0}: \overline{\mathbf{x}} \in \mathcal{Y}[N, \beta]\right) \leq e^{-(\epsilon-\beta) N} \tag{21}
\end{equation*}
$$

Remark 1 The probability $\mathbb{Q}_{0}^{N}(\overline{\mathbf{x}} \in \mathcal{Y}[N, \beta])$ decays exponentially with $N$ only if $\epsilon>\beta$. Thus, uncertainty in the measure manifests itself as a lower bound on the acheivable error probability.

Proof: Fix $0<\eta \leq \epsilon$. Since $\operatorname{err}(\overline{\mathbf{x}})>\epsilon$ we can select $\mathbb{Q}_{1} \in \mathcal{Q}$ such that $\mathbb{Q}_{1}(\mathbf{H}: f(\overline{\mathbf{x}}, \mathbf{H})>$ $0)>\epsilon-\eta$. Let $\mathbf{H}^{0} \sim \mathbb{Q}_{i}, i=1,2$. Since $\rho_{p}\left(\mathbb{Q}_{1}, \mathbb{Q}_{0}\right) \leq \beta$, the Strassen-Dudley Representation Theorem implies that there exists a coupling $\left(\tilde{\mathbf{H}}^{1}, \tilde{\mathbf{H}}^{0}\right)$ of the pair $\left(\mathbf{H}^{1}, \mathbf{H}^{0}\right)$ such that (17) holds, i.e. $\mathbb{P}\left(\left\|\tilde{\mathbf{H}}^{1}-\tilde{\mathbf{H}}^{0}\right\|>\beta\right) \leq \beta$. Let $\left\{\left(\tilde{\mathbf{H}}_{1}^{1}, \tilde{\mathbf{H}}_{1}^{0}\right), \ldots,\left(\tilde{\mathbf{H}}_{N}^{1}, \tilde{\mathbf{H}}_{N}^{0}\right)\right\}$ denote $N$ IID samples of the jointly distributed pair of random vectors ( $\left.\tilde{\mathbf{H}}^{1}, \tilde{\mathbf{H}}^{0}\right)$. Let $\tilde{\mathcal{Y}}[N, \beta]$ denote the set

$$
\begin{equation*}
\tilde{\mathcal{Y}}[N, \beta]=\left\{\mathbf{x}: f(\mathbf{x}, \mathbf{z}) \leq 0, \forall \mathbf{z} \text { s.t. }\left\|\mathbf{z}-\tilde{\mathbf{H}}_{i}^{0}\right\| \leq \beta, i=1, \ldots, N\right\} \tag{22}
\end{equation*}
$$

Since $\mathbf{H}^{0} \stackrel{D}{=} \tilde{\mathbf{H}}^{0}$ and $\overline{\mathbf{x}}$ is fixed, we have that

$$
\begin{equation*}
\mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N}^{0}: \overline{\mathbf{x}} \in \mathcal{Y}[N, \beta]\right)=\mathbb{P}^{N}\left(\tilde{\mathbf{H}}_{1, N}^{0}: \overline{\mathbf{x}} \in \tilde{\mathcal{Y}}[N, \beta]\right) \tag{23}
\end{equation*}
$$

Moreover,

$$
\begin{align*}
\mathbb{P}^{N}\left(\tilde{\mathbf{H}}_{1, N}^{0}: \overline{\mathbf{x}} \in \tilde{\mathcal{Y}}[N, \beta]\right) & =\prod_{i=1}^{N} \mathbb{P}\left(\tilde{\mathbf{H}}_{i}^{0}: f(\overline{\mathbf{x}}, \mathbf{z}) \leq 0, \forall\left\|\mathbf{z}-\tilde{\mathbf{H}}_{i}^{0}\right\| \leq \beta\right) \\
& =\left(\mathbb{P}\left(\tilde{\mathbf{H}}: f(\overline{\mathbf{x}}, \mathbf{z}) \leq 0, \forall\left\|\mathbf{z}-\tilde{\mathbf{H}}^{0}\right\| \leq \beta\right)\right)^{N} \tag{24}
\end{align*}
$$

Each term in (24) can be bounded as follows.

$$
\begin{align*}
& \mathbb{P}\left(f(\overline{\mathbf{x}}, \mathbf{z}) \leq 0, \forall\left\|\mathbf{z}-\tilde{\mathbf{H}}^{0}\right\| \leq \beta\right) \\
& =\mathbb{P}\left(f(\overline{\mathbf{x}}, \mathbf{z}) \leq 0, \forall\left\|\mathbf{z}-\tilde{\mathbf{H}}^{0}\right\| \leq \beta,\left\|\tilde{\mathbf{H}}^{1}-\tilde{\mathbf{H}}^{0}\right\| \leq \beta\right) \\
& +\mathbb{P}\left(f(\overline{\mathbf{x}}, \mathbf{z}) \leq 0, \forall\left\|\mathbf{z}-\tilde{\mathbf{H}}^{0}\right\| \leq \beta,\left\|\tilde{\mathbf{H}}^{1}-\tilde{\mathbf{H}}^{0}\right\|>\beta\right), \\
& \leq \mathbb{P}\left(f\left(\overline{\mathbf{x}}, \tilde{\mathbf{H}}^{1}\right) \leq 0\right)+\mathbb{P}\left(\left\|\tilde{\mathbf{H}}^{1}-\tilde{\mathbf{H}}^{0}\right\|>\beta\right),  \tag{25}\\
& \leq(1-\epsilon+\eta)+\beta, \tag{26}
\end{align*}
$$

where (25) follows from the fact that the probability only increases if one removes restrictions, and (26) follows from the bound (17) and the fact that $\tilde{\mathbf{H}}^{1} \stackrel{\mathrm{D}}{=} \mathbf{H}^{1}$. From (23), (26) and (24), we have

$$
\begin{equation*}
\mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N}^{0}: \overline{\mathbf{x}} \in \mathcal{Y}[N, \beta]\right)=\mathbb{P}^{N}\left(\tilde{\mathbf{H}}_{1, N}^{0}: \overline{\mathbf{x}} \in \tilde{\mathcal{Y}}[N, \beta]\right) \leq(1-\epsilon+\beta+\eta)^{N} \leq e^{-N(\epsilon-\beta-\eta)} \tag{27}
\end{equation*}
$$

Since $\eta \leq \epsilon$ was arbitrary, the result follows.
Note that we only generate samples according to the central measure $\mathbb{Q}_{0}$. The coupling is a construct needed to translate the bound on extremal measure that achieves the supremum in the definition of $\operatorname{err}(\mathbf{x})$ to the measure $\mathbb{Q}_{0}$.

Suppose $\mathbb{Q}(\mathbf{H}:\|\mathbf{H}\|>R)=0$ for all $\mathbb{Q} \in \mathcal{Q}$. Define the set $\mathcal{Y}_{\lambda}[N]$ as follows.

$$
\begin{equation*}
\mathcal{Y}_{\lambda}[N]=\left\{\mathbf{x}: f(\mathbf{x}, \mathbf{z}) \leq 0, \forall \mathbf{z} \text { s.t. }\left\|\mathbf{z}-\mathbf{H}_{i}^{0}\right\| \leq \lambda \beta, i=1, \ldots, N\right\} \tag{28}
\end{equation*}
$$

Then Theorem 5 and Markov's inequality implies the following corollary.
Corollary 3 Fix $\epsilon>0$ and $\overline{\mathbf{x}} \in \mathcal{X}$ with $\operatorname{err}(\overline{\mathbf{x}})>\epsilon$. Suppose $\mathbb{Q}(\mathbf{H}:\|\mathbf{H}\|>R)=0$ for all $\mathbb{Q} \in \mathcal{Q}$. Then

$$
\begin{equation*}
\mathbb{P}\left(\overline{\mathbf{x}} \in \mathcal{Y}_{\lambda}[N]\right) \leq\left(1-\epsilon+\frac{1+2 R}{\lambda}\right)^{N} . \tag{29}
\end{equation*}
$$

Unlike in Lemma 6, here we have a parameter $\lambda$ that can be controlled to achieve any desired probability of error $\epsilon$.

Next, we establish a robust analog of Lemma 2. We show that if the VC dimension $d_{f}$ of the function $f(\cdot, \cdot)$ is finite, $\rho_{p}\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \beta$, and the number of samples $N=\mathcal{O}\left(\frac{d_{f}}{\epsilon-\beta}\right)$ (a precise bound is given in Lemma 7),

$$
\mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N}^{0}: \mathcal{Y}[N, \beta] \subseteq \mathcal{X}_{\epsilon}(\mathbb{Q})\right) \geq 1-\delta .
$$

This result should be interpreted as follows. The distribution of the parameters $\mathbf{H}$ is uncertain and is only known to lie in the uncertainty set $\mathcal{Q}=\left\{\mathbb{Q}: \rho_{p}\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \beta\right\}$ and we want to characterize the set of decisions $\mathbf{x}$ that satisfy $\mathbb{Q}(\mathbf{H}: f(\mathbf{x}, \mathbf{H})>0) \leq \epsilon$ no matter which probability measure $\mathbb{Q}$ is selected from the uncertainty set $\mathcal{Q}$. The bound above shows that for $N=\mathcal{O}\left(\frac{d_{f}}{\epsilon-\beta}\right)$ the set $\mathcal{Y}[N, \beta]$ is a good approximation for $\mathcal{X}_{\epsilon}(\mathbb{Q})$ for any fixed $\mathbb{Q}$ with high probability.

Lemma 7 Fix $\delta>0, \epsilon>\beta$ and $\mathbb{Q}_{1} \in \mathcal{Q}$. Suppose the VC dimension $d_{f}$ of the function $f(\cdot, \cdot)$ is finite and $\beta+2^{-\beta / 2}<1$. Then $\mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N}^{0}: \mathcal{Y}[N, \beta] \nsubseteq \mathcal{X}_{\epsilon}\left(\mathbb{Q}_{1}\right)\right) \leq \delta$, for all $N$ satisfying

$$
\begin{aligned}
& N \geq \max \left\{\begin{array}{cc}
d, \quad \frac{8}{\epsilon}, \quad \frac{2 d_{f}}{e(1-\beta)} \ln \left(\frac{e}{1-\beta}\right)+\frac{2}{1-\beta} \ln \left(\frac{e}{(e-1) \delta}\right)+1, ~
\end{array}\right. \\
& \left.\frac{4 d_{f}}{\epsilon-\mu} \log \left(\frac{12}{\epsilon-\mu}\right)+\frac{4}{\epsilon-\mu} \log \left(\frac{2}{\delta(1-\beta)}\right)\right\},
\end{aligned}
$$

where $\mu=2\left(\frac{\epsilon}{2}+\log \left(\beta+2^{-\epsilon / 2}\right)\right)$.
Remark 2 Since $\beta=0$ implies $\mu=0$, we recover the non-robust result in Lemma 2 when $\beta=0$.
Proof: Since the measure $\mathbb{Q}_{1} \in \mathcal{Q}$ is fixed, we will abbreviate $\mathcal{X}_{\epsilon}\left(\mathbb{Q}_{1}\right)$ by $\mathcal{X}_{\epsilon}$. Let $\mathcal{X}_{\epsilon}^{c}$ denote the complement of the set $\mathcal{X}_{\epsilon}$. As in the proof of Lemma 6, let $\mathbf{H}^{i} \sim \mathbb{Q}_{i}, i=1,2$ and let $\left(\tilde{\mathbf{H}}^{1}, \tilde{\mathbf{H}}^{0}\right)$ denote a coupling of the pair $\left(\mathbf{H}^{1}, \mathbf{H}^{0}\right)$ such that (17) holds, i.e. $\mathbb{P}\left(\left\|\tilde{\mathbf{H}}^{1}-\tilde{\mathbf{H}}^{0}\right\|>\beta\right) \leq \beta$. Let $\left\{\left(\tilde{\mathbf{H}}_{1}^{1}, \tilde{\mathbf{H}}_{1}^{0}\right), \ldots,\left(\tilde{\mathbf{H}}_{N}^{1}, \tilde{\mathbf{H}}_{N}^{0}\right)\right\}$ denote $N$ IID samples of the jointly distributed pair of random vectors $\left(\tilde{\mathbf{H}}^{1}, \tilde{\mathbf{H}}^{0}\right)$. Then

$$
\begin{aligned}
\mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N}^{0}: \mathcal{Y}[N, \beta] \nsubseteq \mathcal{X}_{\epsilon}\right) & =\mathbb{P}^{N}\left(\mathbf{H}_{1, N}^{0}: \tilde{\mathcal{Y}}[N, \beta] \cap \mathcal{X}_{\epsilon}^{c} \neq \emptyset\right), \\
& =\sum_{j=0}^{N} \mathbb{P}^{N}\left(\tilde{\mathcal{Y}}[N, \beta] \cap \mathcal{X}_{\epsilon}^{c} \neq \emptyset,|\mathcal{I}|=j\right),
\end{aligned}
$$

where $\mathcal{I}=\left\{i \in\{1, \ldots, N\}:\left\|\tilde{\mathbf{H}}_{i}^{1}-\tilde{\mathbf{H}}_{i}^{0}\right\| \leq \beta\right\}$. For a set $\mathcal{I} \subset\{1, \ldots, N\}$ let $\mathcal{A}(\mathcal{I})$ denote the event

$$
\mathcal{A}(\mathcal{I})=\left\{\left(\tilde{\mathbf{H}}_{i}^{0}, \tilde{\mathbf{H}}_{i}^{1}\right)_{i=1, \ldots, N}:\left\|\tilde{\mathbf{H}}_{k}^{1}-\tilde{\mathbf{H}}_{k}^{0}\right\| \leq \beta, \forall k \in \mathcal{I},\left\|\tilde{\mathbf{H}}_{k}^{1}-\tilde{\mathbf{H}}_{k}^{0}\right\| \leq \beta, \forall k \notin \mathcal{I}\right\}
$$

and let $\mathcal{Y}[\mathcal{I}, \beta]=\left\{\mathbf{x}: f(\mathbf{x}, \mathbf{z}) \leq 0, \forall \mathbf{z}\right.$ s.t. $\left.\left\|\mathbf{z}-\tilde{\mathbf{H}}_{i}^{0}\right\| \leq \beta, i \in \mathcal{I}_{j}\right\}$. Fix $\mathcal{I}_{1}, \mathcal{I}_{2} \subseteq\{1, \ldots, N\}$ with $\left|\mathcal{I}_{1}\right|=\left|\mathcal{I}_{2}\right|$. Since $\left\{\left(\tilde{\mathbf{H}}_{i}{ }^{1}, \tilde{\mathbf{H}}_{i}{ }^{0}\right)\right\}, i=1, \ldots, N$ are IID, it is clear that

$$
\begin{align*}
& \mathbb{P}^{N}\left(\left(\tilde{\mathbf{H}}_{i}^{0}, \tilde{\mathbf{H}}_{i}^{1}\right)_{i=1, \ldots, N}: \tilde{\mathcal{Y}}[N, \beta] \cap \mathcal{X}_{\epsilon}^{c} \neq \emptyset, \mathcal{A}\left(\mathcal{I}_{1}\right)\right) \\
& \quad=\mathbb{P}^{N}\left(\left(\tilde{\mathbf{H}}_{i}^{0}, \tilde{\mathbf{H}}_{i}^{1}\right)_{i=1, \ldots, N}: \tilde{\mathcal{Y}}[N, \beta] \cap \mathcal{X}_{\epsilon}^{c} \neq \emptyset, \mathcal{A}\left(\mathcal{I}_{2}\right)\right) . \tag{30}
\end{align*}
$$

Set $\mathcal{I}_{0}=\emptyset$, and $\mathcal{I}_{j}=\{1, \ldots, j\}, j=1, \ldots, N$. Since there are $\binom{N}{j}$ possible selections for the set $\mathcal{I}_{j}$ of cardinality $j$, (30) implies that

$$
\begin{align*}
\sum_{j=0}^{N} & \mathbb{P}^{N}\left(\left(\tilde{\mathbf{H}}_{i}^{0}, \tilde{\mathbf{H}}_{i}^{1}\right)_{i=1, \ldots, N}: \tilde{\mathcal{Y}}[N, \beta] \cap \mathcal{X}_{\epsilon}^{c} \neq \emptyset,|\mathcal{I}|=j\right) \\
& =\sum_{j=0}^{N}\binom{N}{j} \mathbb{P}^{N}\left(\left(\tilde{\mathbf{H}}_{i}^{0}, \tilde{\mathbf{H}}_{i}^{1}\right)_{i=1, \ldots, N}: \tilde{\mathcal{Y}}[N, \beta] \cap \mathcal{X}_{\epsilon}^{c}, \mathcal{A}\left(\mathcal{I}_{j}\right)\right), \\
& \leq \sum_{j=0}^{N}\binom{N}{j} \mathbb{P}^{N}\left(\left(\tilde{\mathbf{H}}_{i}^{0}, \tilde{\mathbf{H}}_{i}^{1}\right)_{i=1, \ldots, N}: \mathcal{X}_{\epsilon}^{c} \cap \mathcal{Y}\left[\mathcal{I}_{j}, \beta\right] \neq \emptyset, \mathcal{A}\left(\mathcal{I}_{j}\right)\right)  \tag{31}\\
& =\sum_{j=0}^{N}\binom{N}{j} \mathbb{P}^{j}\left(\left(\tilde{\mathbf{H}}_{k}^{0}, \tilde{\mathbf{H}}_{k}^{1}\right)_{k \in \mathcal{I}_{j}}: \mathcal{X}_{\epsilon}^{c} \cap \mathcal{Y}\left[\mathcal{I}_{j}, \beta\right] \neq \emptyset, \mathcal{A}\left(\mathcal{I}_{j}\right)\right) \\
& \leq \sum_{j=0}^{N}\binom{N}{j} \beta^{N-j} \mathbb{P}^{j}\left(\left(\tilde{\mathbf{H}}_{i}^{0}, \tilde{\mathbf{H}}_{i}^{1}\right)_{i=1, \ldots, j}: \mathcal{X}_{\epsilon}^{c} \cap \mathcal{Y}\left[\mathcal{I}_{j}, \beta\right] \neq \emptyset, \mathcal{A}\left(\mathcal{I}_{j}\right)\right)  \tag{32}\\
& \leq \sum_{j=0}^{N}\binom{N}{j} \beta^{N-j} \mathbb{P}^{j}\left(\tilde{\mathbf{H}}_{1, j}^{1}: \exists \mathbf{x} \in \mathcal{X}_{\epsilon}^{c} \text { s.t. } f\left(\mathbf{x}, \tilde{\mathbf{H}}_{k}^{1}\right) \leq 0, \forall k \in \mathcal{I}_{j}\right), \tag{33}
\end{align*}
$$

where (31) and (34) follows from the fact that the probability only increases if one removes restrictions, (32) follows from the fact that $\left\{\left(\tilde{\mathbf{H}}_{i}{ }^{1}, \tilde{\mathbf{H}}_{i}{ }^{0}\right)\right\}, i=1, \ldots, N$ are IID, and (33) follows from the bound (17). Note that the bound (34) only involves the random vector $\mathbf{H}^{1}$, or equivalently the (unknown) true measure $\mathbb{Q}_{1}$. Thus, once again we have used coupling to translate a bound in terms of the central measure $\mathbb{Q}_{0}$ to one involving the measure $\mathbb{Q}_{1}$. We do not need coupling beyond this stage of the proof. In the rest of this proof we bound (34) using Lemma 2 applied to the (unknown)
measure $\mathbb{Q}_{1}$. Let $N_{0}=\left\lfloor\frac{8}{\epsilon}\right\rfloor$. Then

$$
\begin{align*}
& \sum_{j=0}^{N}\binom{N}{j} \beta^{N-j} \mathbb{P}^{j}\left(\tilde{\mathbf{H}}_{1, j}^{1}: \exists \mathbf{x} \in \mathcal{X}_{\epsilon}^{c} \text { s.t. } f\left(\mathbf{x}, \tilde{\mathbf{H}}_{k}^{1}\right) \leq 0, \forall k \in \mathcal{I}_{j}\right) \\
&= \sum_{j=0}^{N_{0}}\binom{N}{j} \beta^{N-j} \mathbb{P}^{j}\left(\tilde{\mathbf{H}}_{1, j}^{1}: \exists \mathbf{x} \in \mathcal{X}_{\epsilon}^{c} \text { s.t. } f\left(\mathbf{x}, \tilde{\mathbf{H}}_{k}^{1}\right) \leq 0, \forall k \in \mathcal{I}_{j}\right) \\
&+\sum_{j=N_{0}+1}^{N}\binom{N}{j} \beta^{N-j} \mathbb{P}^{j}\left(\tilde{\mathbf{H}}_{1, j}^{1}: \exists \mathbf{x} \in \mathcal{X}_{\epsilon}^{c} \text { s.t. } f\left(\mathbf{x}, \tilde{\mathbf{H}}_{k}^{1}\right) \leq 0, \forall k \in \mathcal{I}_{j}\right), \\
& \leq \sum_{j=0}^{N_{0}}\binom{N}{j} \beta^{N-j}+\sum_{j=N_{0}+1}^{N}\binom{N}{j} \beta^{N-j} \mathbb{P}^{j}\left(\tilde{\mathbf{H}}_{1, j}^{1}: \exists \mathbf{x} \in \mathcal{X}_{\epsilon}^{c} \text { s.t. } f\left(\mathbf{x}, \tilde{\mathbf{H}}_{k}^{1}\right) \leq 0, \forall k \in \mathcal{I}_{j}\right) \\
& \leq \sum_{j=0}^{N_{0}}\binom{N}{j} \beta^{N-j}+\sum_{j=N_{0}+1}^{N}\binom{N}{j} \beta^{N-j}\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon j / 2}, \tag{35}
\end{align*}
$$

where (35) follows from Lemma 2 and the bound (13). The rest of this proof is tedious algebra to prove a "nice" bound on (35).

$$
\begin{align*}
& \sum_{j=0}^{N_{0}}\binom{N}{j} \beta^{N-j}+\sum_{j=N_{0}+1}^{N}\binom{N}{j} \beta^{N-j}\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon j / 2} \\
& \quad=\underbrace{\sum_{j=0}^{N_{0}}\binom{N}{j} \beta^{N-j}\left(1-\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon j / 2}\right)}_{\tau_{1}}+\underbrace{\sum_{j=0}^{N}\binom{N}{j} \beta^{N-j}\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon j / 2}}_{\tau_{2}} \tag{36}
\end{align*}
$$

To complete the proof we show that if $N$ is large enough the terms $\tau_{1}$ and $\tau_{2}$ are bounded by $\tau_{1} \leq \delta \beta$ and $\tau_{2} \leq \delta(1-\beta)$, which implies that $\tau_{1}+\tau_{2} \leq \delta$. We can bound $\tau_{1}$ as follows. Let $d_{0}=\left\lfloor\frac{d_{f}}{e}\right\rfloor$ where $e$ is the base of natural logarithm. Then

$$
\tau_{1}=\sum_{j=0}^{d_{0}}\binom{N}{j} \beta^{N-j}\left(1-\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon j / 2}\right)+\sum_{j=d_{0}+1}^{N_{0}}\binom{N}{j} \beta^{N-j}\left(1-\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon j / 2}\right)
$$

Note that for $\frac{d_{f}}{e} \leq d_{0}+1 \leq j \leq N_{0} \leq \frac{8}{\epsilon}$. Thus, we have

$$
\begin{aligned}
1-\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon j / 2} & \leq 1-\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon N_{0} / 2} \\
& \leq 1-\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon \frac{8}{2 \epsilon}} \\
& =1-\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{-3} \\
& \leq 1-\left(\frac{2 e d_{f}}{d_{f} e}\right)^{d_{f}} 2^{-3} \\
& =1-2^{d_{f}-3} \leq 0
\end{aligned}
$$

The last inequality follows from the assumption that $d_{f}>3$. Therefore,

$$
\begin{align*}
\tau_{1} & \leq \sum_{j=0}^{d_{0}}\binom{N}{j} \beta^{N-j}\left(1-\left(\frac{2 e j}{d}\right)^{d} 2^{1-\epsilon j / 2}\right), \\
& \leq \sum_{j=0}^{d_{0}}\binom{N}{j} \beta^{N-j}, \\
& \leq \frac{\beta N}{N-d_{0}} \sum_{j=0}^{d_{0}}\binom{N-1}{j} \beta^{N-1-j}, \\
& \leq \frac{N \beta(1-\beta)^{-d_{0}}}{N-d_{0}} \sum_{j=0}^{d_{0}}\binom{N-1}{j} \beta^{N-1-j}(1-\beta)^{j}, \\
& =\left(\frac{N \beta(1-\beta)^{-d_{0}}}{N-d_{0}}\right) P\left(1-\beta, N-1, d_{0}\right), \tag{37}
\end{align*}
$$

where $P(p, N, s)$ denotes the probability of at most $s$ successes in $N$ IID Bernoulli trials, each with a success probability $p$. Let $\theta=1-\frac{d_{0}}{(N-1)(1-\beta)}$. Then, Chernoff bound implies that

$$
\tau_{1} \leq \frac{N \beta(1-\beta)^{-d_{0}}}{N-d_{0}} \exp \left\{-\frac{(N-1)(1-\beta)}{2}+d_{0}\right\}
$$

For $N \geq d_{f} \geq e d_{0}$ we have $\frac{N}{N-d_{0}} \leq \frac{e}{e-1}$. Therefore,

$$
\begin{equation*}
\tau_{1} \leq \frac{e \beta(1-\beta)^{-d_{f} / \epsilon}}{e-1} \exp \left\{-\frac{(N-1)(1-\beta)}{2}+\frac{d}{\epsilon}\right\} \tag{38}
\end{equation*}
$$

Thus, $\tau_{1} \leq \delta \beta$ for all

$$
\begin{equation*}
N \geq \frac{2 d_{f}}{e(1-\beta)} \ln \left(\frac{e}{1-\beta}\right)+\frac{2}{1-\beta} \ln \left(\frac{e}{(e-1) \delta}\right)+1 \tag{39}
\end{equation*}
$$

Next, we bound $\tau_{2}$ as follows.

$$
\begin{align*}
\tau_{2} & =\sum_{j=0}^{N}\binom{N}{j} \beta^{N-j}\left(\frac{2 e j}{d_{f}}\right)^{d_{f}} 2^{1-\epsilon j / 2} \\
& =2\left(\frac{2 e}{d_{f}}\right)^{d_{f}} \sum_{j=0}^{N}\binom{N}{j} j^{d_{f}} \beta^{N-j} 2^{-\epsilon j / 2}, \\
& \leq 2\left(\frac{2 e}{d_{f}}\right)^{d_{f}} N^{d_{f}} \sum_{j=0}^{N}\binom{N}{j} \beta^{N-j} 2^{-\epsilon j / 2}, \\
& \leq 2\left(\frac{2 e}{d_{f}}\right)^{d_{f}} N^{d_{f}}\left(\beta+2^{-\epsilon / 2}\right)^{N} \tag{40}
\end{align*}
$$

Since $\beta+2^{-\epsilon / 2} \leq \beta+2^{-\beta / 2}<1, \mu=2\left(\frac{\epsilon}{2}+\log \left(\beta+2^{-\epsilon / 2}\right)\right)$ is well defined. Then an analysis similar to the one given in [9] (see also [1] pg. 95 for details) shows that

$$
\tau_{2} \leq 2\left(\frac{2 e}{d_{f}}\right)^{d_{f}} N^{d_{f}}\left(2^{-(\epsilon-\mu) / 2}\right)^{N}
$$

Thus, $\tau_{2} \leq(1-\delta) \beta$ for all

$$
\begin{equation*}
N \geq \frac{4 d_{f}}{\epsilon-\mu} \log \left(\frac{12}{\epsilon-\mu}\right)+\frac{4}{\epsilon-\mu} \log \left(\frac{2}{\delta(1-\beta)}\right) \tag{41}
\end{equation*}
$$

The result follows from (36), (39), and (41).
The last result in this section is the robust analog of Lemma 4.
Lemma 8 Fix $\epsilon>0$. Let $\widehat{\mathbf{x}}$ denote the optimal solution of the robust sampled problem (8). Then $\mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N}^{0}: \widehat{\mathbf{x}} \notin \overline{\mathcal{X}}_{\epsilon}\right) \leq\left(\frac{e N}{n}\right)^{n} e^{-(\epsilon-\beta)(N-n)}$.

Proof: The robust chance constrained problem (8) has constraints of the form

$$
f(\mathbf{x}, \mathbf{z}) \leq 0, \quad\left\|\mathbf{z}-\mathbf{H}_{i}^{0}\right\| \leq \beta, \quad i=1, \ldots, N .
$$

Suppose a constraint of the form $f(\mathbf{x}, \overline{\mathbf{z}}) \leq 0$ is a support constraint for the robust chance constrained problem (8). We will associate this support constraint with $k=\operatorname{argmin}\left\{i:\left\|\overline{\mathbf{z}}-\mathbf{H}_{i}^{0}\right\| \leq \beta\right\}$.

Let $\mathcal{I} \subseteq\{1, \ldots, N\}$ with $|I| \leq n$ and let

$$
\mathcal{H}_{\mathcal{I}}^{N}=\left\{\left(\mathbf{h}_{1}, \ldots, \mathbf{h}_{N}\right): \text { all support constraints are associated with some } i \in \mathcal{I}\right\} .
$$

Then Theorem 3 implies $\mathcal{H}^{N}=\cup_{\{\mathcal{I} \subseteq\{1, \ldots, N\}:|\mathcal{I}| \leq n\}} \mathcal{H}_{\mathcal{I}}^{N}$. Thus,

$$
\begin{aligned}
\mathbb{Q}_{0}^{N} & \left(\mathbf{H}_{1, N}: \widehat{\mathbf{x}} \notin \overline{\mathcal{X}}_{\epsilon}\right) \\
& \sum_{\{\mathcal{I} \subseteq\{1, \ldots, N\}:|\mathcal{I}| \leq n\}} \mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N} \in \mathcal{H}_{\mathcal{I}}^{N}: \widehat{\mathbf{x}}_{\mathcal{I}} \notin \overline{\mathcal{X}}_{\epsilon}\right) \\
& =\sum_{\{\mathcal{I} \subseteq\{1, \ldots, N\}:|\mathcal{I}|=n\}}\left(\mathbb{Q}_{0}^{n}\left(\mathbf{H}_{i \in I}: \widehat{\mathbf{x}}_{\mathcal{I}} \notin \overline{\mathcal{X}}_{\epsilon}\right) \prod_{i \notin \mathcal{I}} \mathbb{Q}_{0}\left(\mathbf{H}_{i}: f\left(\widehat{\mathbf{x}}_{\mathcal{I}}, \mathbf{z}\right) \leq 0, \forall \forall\left\|\mathbf{z}-\mathbf{H}_{i}\right\| \leq \beta \mid \mathcal{A}_{\mathcal{I}}\right)\right),
\end{aligned}
$$

where $\widehat{\mathbf{x}}_{\mathcal{I}}$ denotes the optimal solution of the robust sampled problem (5) with only the robust constraints corresponding to the samples $i \in \mathcal{I}$ present, $\mathcal{A}_{\mathcal{I}}$ is the event $\mathcal{A}_{\mathcal{I}}=\left\{\mathbf{H}_{i \in I}: \widehat{\mathbf{x}}_{\mathcal{I}} \notin \overline{\mathcal{X}}_{\epsilon}\right\}$ and each term in the sum can be written as the product because $\mathbf{H}_{1, N}^{0}$ are IID samples. Lemma 1 implies that $\mathbb{Q}_{0}\left(\mathbf{H}: f\left(\widehat{\mathbf{x}}_{\mathcal{I}}, \mathbf{z}\right) \leq 0, \forall\|\mathbf{z}-\mathbf{H}\| \leq \beta \mid \mathcal{A}_{\mathcal{I}}\right) \leq e^{-(\epsilon-\beta)}$. Thus,

$$
\begin{aligned}
\mathbb{Q}_{0}^{N}\left(\mathbf{H}_{1, N}: \widehat{\mathbf{x}} \notin \mathcal{X}_{\epsilon}\right) & \leq e^{-(\epsilon-\beta)(N-n)} \sum_{\{\mathcal{I} \subseteq\{1, \ldots, N\}:|\mathcal{I}|=n\}} \mathbb{Q}_{0}^{n}\left(\mathbf{H}_{i \in I}: \widehat{\mathbf{x}}_{\mathcal{I}} \notin \overline{\mathcal{X}}_{\epsilon}\right), \\
& \leq e^{-(\epsilon-\beta)(N-n)}\left(\sum_{k=1}^{n}\binom{N}{k}\right) \leq\left(\frac{e N}{n}\right)^{n} e^{-(\epsilon-\beta)(N-n)},
\end{aligned}
$$

where the last inequality follows from the bound (13).

## 6 Tractability of the robust sampled problem

In Section 1 we introduced the robust sampled problem (8) as an approximation for the ambiguous chance constrained problem (6) and in Section 5 we established bounds on the number of samples $N$ required to approximate the robust feasible set $\overline{\mathcal{X}}_{\epsilon}$ and also on the number of samples required to only ensure that the optimal solution $\widehat{\mathbf{x}}$ of the robust problem (8) is feasible for (6). All along we have implicitly assumed that the robust sampled problem (8) is efficiently solvable. In this section, we characterize the functions $f(\cdot, \cdot)$, the probability metric $\rho$ and the norm $\|\cdot\|$ on the parameter space $\mathcal{H}$ for which the robust sampled problem (8) is tractable both in theory and in practice. The results in this section are motivated by [7].

We restrict attention to the following two classes of constraint functions.
(a) Affine functions: $\mathcal{X}=\mathbf{R}^{n}, \mathcal{H}=\mathbf{R}^{n+1}$, and $f\left(\mathbf{x},\left(h_{0}, \mathbf{h}\right)\right)=h_{0}+\mathbf{h}^{T} \mathbf{x}$.
(b) Second-order cone functions: $\mathbf{x} \in \mathbf{R}^{n}, \mathcal{H}=\left\{\mathbf{h}=(\mathbf{A}, \mathbf{b}, \mathbf{u}, v): \mathbf{A} \in \mathbf{R}^{p \times n}, \mathbf{b} \in \mathbf{R}^{p}, \mathbf{u} \in \mathbf{R}^{n}, v \in \mathbf{R}\right\}$, and $f(\mathbf{x}, \mathbf{h})=\sqrt{(\mathbf{A x}+\mathbf{b})^{T}(\mathbf{A x}+\mathbf{b})}-\mathbf{u}^{T} \mathbf{x}-v$.

The uncertainty set $\mathcal{Q}$ considered in this paper is given by $\mathcal{Q}=\left\{\mathbb{Q}: \rho_{p}\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \beta\right\}$ where $\rho_{p}$ denotes the Prohorov metric. Since the Prohorov metric is defined in terms of the norm $\|\cdot\|$ on the space $\mathcal{H}$, we first select this norm. We restrict attention to norms that satisfy

$$
\begin{equation*}
\|\mathbf{u}\|=\||\mathbf{u}|\|, \tag{42}
\end{equation*}
$$

where $|\mathbf{u}|$ denotes the vector obtained by taking the absolute value of each of the components. For a given norm $\|\cdot\|$, the constant $\beta$ defining the uncertainty set $\mathcal{Q}$ is set by the desired level of confidence. Note that $\beta$ can also be set in terms of any distance measure that is an upper bound for the Prohorov metric. See Section 4 for details.

First we consider the case of affine constraint functions $f(\mathbf{x}, \mathbf{h})=h_{0}+\mathbf{h}^{T} \mathbf{x}$. Let $\mathbf{e}_{j}, j=$ $1, \ldots, n+1$ denote the $j$-th basis vector in $\mathbf{R}^{n+1}$. Define $\mathcal{U}(\mathbf{h})=\left\{\mathbf{z}: \mathbf{z}=\mathbf{h}+\sum_{i=1}^{n+1} w_{j} \mathbf{e}_{j},\|\mathbf{w}\| \leq \beta\right\}$. Then the robust sampled problem (8) is given by

$$
\begin{align*}
\min & \mathbf{c}^{T} \mathbf{x} \\
\text { s.t. } & z_{0}+\mathbf{z}^{T} \mathbf{x} \leq 0, \quad \forall \mathbf{z} \in \mathcal{U}\left(\mathbf{H}_{i}^{0}\right), \quad i=1, \ldots, N,  \tag{43}\\
& \mathbf{x} \in \mathcal{X}
\end{align*}
$$

Results in [7] show that (43) can be reformulated as follows.

$$
\begin{array}{cll}
\min & \mathbf{c}^{T} \mathbf{x} & \\
\text { s.t. } & f\left(\mathbf{x}, \mathbf{H}_{i}^{0}\right) \leq-\beta y_{i}, & i=1, \ldots, N, \\
& \left|x_{j}\right| \leq t_{j}^{i}, & j=1, \ldots, n, \quad i=1, \ldots, N, \\
& 1 \leq t_{n+1}^{i}, & i=1, \ldots, N,  \tag{44}\\
& \left\|\mathbf{t}^{i}\right\|_{*} \leq y_{i}, & i=1, \ldots, N, \\
& \mathbf{y} \in \mathbf{R}^{N}, \mathbf{t}^{i} \in \mathbf{R}^{n+1}, & i=1, \ldots, N, \\
& \mathbf{x} \in \mathcal{X}
\end{array}
$$

where $\|\mathbf{s}\|_{*}=\max _{\{\|\mathbf{r}\| \leq 1\}}\left\{\mathbf{s}^{T} \mathbf{r}\right\}$ denotes the dual norm of $\|\cdot\|$. For the $\mathcal{L}_{1}$ or $\mathcal{L}_{\infty}$ norms (44) reduces to a linear program; whereas when the norm $\|\cdot\|$ is an $\mathcal{L}_{p}$-norm, $p \neq\{1, \infty\}$, (44) is equivalent to a second-order cone program.

Next, consider the same of the second-order cone constraints. Let $\mathbf{e}_{j} \in \mathbf{R}^{(p+1)(n+1)}$ denote the $j$-th standard basis vector in $\mathbf{R}^{(p+1)(n+1)}$. For $\mathbf{A}=\left[\mathbf{a}_{1}, \ldots, \mathbf{a}_{n}\right] \in \mathbf{R}^{p \times n}$ let

$$
\operatorname{vec}(\mathbf{A})=\left[\begin{array}{lll}
\mathbf{a}_{1}^{T} & \cdots & \mathbf{a}_{n}^{T}
\end{array}\right]^{T} \in \mathbf{R}^{p n}
$$

and, for $\mathbf{h} \in \mathbf{R}^{(p+1)(n+1)}$, define $\mathcal{U}(\mathbf{h})=\left\{\mathbf{z}: \mathbf{z}=\mathbf{h}+\sum_{j=1}^{(p+1)(n+1)} w_{j} \mathbf{e}_{j},\|\mathbf{w}\| \leq \beta\right\}$. It is shown in [7] that any feasible solution to the problem (45) below is also feasible for the robust sampled problem (8).

$$
\begin{array}{cll}
\min & \mathbf{c}^{T} \mathbf{x} & \\
\text { s.t. } & f\left(\mathbf{x}, \mathbf{H}_{i}^{0}\right) \leq-\beta y_{i}, & i=1, \ldots, N, \\
& g_{j}^{i}(\mathbf{x}) \leq t_{j}^{i}, & i=1, \ldots,(p+1)(n+1), \quad i=1, \ldots, N, \\
& \left\|\mathbf{t}^{i}\right\|_{*} \leq y_{i}, & \\
& \mathbf{y} \in \mathbf{R}^{N}, \mathbf{t}^{i} \in \mathbf{R}^{(p+1)(n+1)}, & i=1, \ldots, N, \\
& \mathbf{x} \in \mathcal{X}, &
\end{array}
$$

where

$$
g_{j}^{i}(\mathbf{x})= \begin{cases}\left|x_{l}\right| & j=p(l-1)+k, \quad k=1, \ldots, p, \quad l=1, \ldots, n \\ 1 & j=p n+k, \quad k=1, \ldots, p \\ \left|x_{l}\right| & j=(p+1) n+l, \quad l=1, \ldots, n \\ 1 & j=(p+1)(n+1),\end{cases}
$$

The problem (45) is a second-order cone program for all $\mathcal{L}_{p}$ norms.

## 7 Conclusion

In this paper we extend the sample complexity results known for chance constrained problems to ambiguous chance constrained problems where the uncertainty set is given by a ball defined in terms of the Prohorov metric. We approximate the ambiguous chance constrained problem by a robust sampled problem where each constraint is a robust constraint centered at a sample drawn according to the center of the uncertainty set defining the ambiguous chance constrained problem. The main contribution of this paper is to show that the robust sampled problem is a good approximation for the ambiguous chance constrained problem with high probability. Our extensions are based on the Strassen-Dudley Representation Theorem that states that when the distributions of two random variables are close in the Prohorov metric one can construct a coupling of the random variables such that the samples are close with high probability. Coupling is just a construct needed to prove the results; it is never used in computing the solution to the robust sampled problem.

The results in this paper should be viewed as a first step towards solving ambiguous chance constrained problems. Several issues still remain unresolved. We only consider uncertainty sets that are norm balls defined in terms of the Prohorov metric. One could consider "tiling" a more general uncertainty set by norm balls of a given radius and construct a robust sampled problem by drawing samples according to the centers of the balls (a simplified version of this idea appears in [11]). Since the constant $\epsilon$ that controls the violation probability in the ambiguous chance constrained problem has to be greater that the radius $\beta$ of the norm ball, such an approach is attractive even when the uncertainty is a norm ball. However, it is not clear how to efficiently select the centers of the balls to "tile" the uncertainty set.

In Section 4 we introduce several probability metrics and show that uncertainty set $\mathcal{Q}=\{\mathbb{Q}$ : $\left.\rho\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \delta\right\}$ can be conservatively approximated by a uncertainty set $\tilde{\mathcal{Q}}=\left\{\mathbb{Q}: \rho_{p}\left(\mathbb{Q}, \mathbb{Q}_{0}\right) \leq \delta(\beta)\right\}$ defined in terms of the Prohorov metric. However, we have no way of measuring the "blow-up" of the uncertainty set that occurs in changing the metrics. This issue can be resolved by either establishing tight lower bounds on the Prohorov metric or by constructing Representation results for the other metrics. Ideally one would like to get logarithmically separated upper and lower bounds as in [32, 33].

Finally, there is the issue of proving worst-case lower bounds on the number of samples required to learn the solution of an ambiguous chance constrained problem, i.e. a refinement of Lemma 3 that accounts for ambiguity in the measure.

## References

[1] M. Anthony and N. Biggs. Computational Learning Theory. Cambridge University Press, 1992.
[2] P. Artzner, F. Delbaen, J-P. Eber, and D. Heath. Coherent risk measures. Math. Finance, 9(203-228), 1999.
[3] A. Ben-Tal and A. Nemirovski. Robust truss topology design via semidefinite programming. SIAM J. Optim., 7(4):991-1016, 1997.
[4] A. Ben-Tal and A. Nemirovski. Robust convex optimization. Math. Oper. Res., 23(4):769-805, 1998.
[5] A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. Oper. Res. Lett., 25(1):1-13, 1999.
[6] A. Ben-Tal and A. Nemirovski. Lectures on Modern Convex Optimization. SIAM, Philadelphia, PA, 2001.
[7] D. Bertsimas and M. Sim. Robust conic optimization. Under review in Math. Prog., May 2004.
[8] J. R. Birge and J. R. B. Wets. Designing approximation schemes for stochastic optimization problems, in particular for stochastic programs with recourse. Math. Prog. Study, 27(54-102), 1986.
[9] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth. Learnability and the VapnikChervonenkis dimension. Journal of the ACM, 36(4):929-965, Oct 1989.
[10] G. Calafiore and M. C. Campi. Uncertain convex programs: Randomized solutions and confidence levels. To appear in Math. Prog., 2003.
[11] G. Calafiore and M. C. Campi. Decision making in an uncertain environment: the scenariobased optimization approach. Working paper, 2004.
[12] A. Charnes and W. W. Cooper. Uncertain convex programs: randomized solutions and confidence levels. Mgmt. Sc., 6:73-79, 1959.
[13] Z. Chen and L.G. Epstein. Ambiguity, risk and asset returns in continuous time. Mimeo, 2000.
[14] D. P. de Farias and B. Van Roy. On constraint sampling in the linear programming approach to approximate dynamic programming. To appear in Math. Oper. Res., 2001.
[15] R. M. Dudley. Distance of probability measures and random variables. Ann. Math. Stat., 39:1563-1572, 1968.
[16] J. Dupačová. The minimax approach to stochastic program and illustrative application. Stochastics, 20:73-88, 1987.
[17] J. Dupačová. Stochastic programming: minimax approach. In Encyclopedia of Optimization. Kluwer, 2001.
[18] L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. SIAM J. Matrix Anal. Appl., 18(4):1035-1064, 1997.
[19] L. El Ghaoui and A. Nilim. The curse of uncertainty in dynamic programming and how to fix it. To appear in Oper. Res.. UC Berkeley Tech Report UCB-ERL-M02/31, November 2002.
[20] L. El Ghaoui, F. Oustry, and H. Lebret. Robust solutions to uncertain semidefinite programs. SIAM J. Optim., 9(1):33-52, 1998.
[21] L. G. Epstein and M. Schneider. Recursive multiple priors. Technical Report 485, Rochester Center for Economic Research, August 2001. Available at http://rcer.econ.rochester.edu. To appear in J. Econ. Theory.
[22] L. G. Epstein and M. Schneider. Learning under Ambiguity. Technical Report 497, Rochester Center for Economic Research, October 2002. Available at http://rcer.econ.rochester.edu.
[23] H. Föllmer and A. Schied. Convex measures of risk and trading constraints. Fin. and Stoch., 6:429-447, 2002.
[24] A. L. Gibbs and F. E. Su. On choosing and bounding probability metrics. Intl. Stat. Rev., 7(3):419-435, 2002.
[25] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. J. Math. Econ., 18(2):141-153, 1989.
[26] D. Goldfarb and Iyengar G. Robust portfolio selection problems. Math. Oper. Res., 28(1):1-38, 2003.
[27] L. P. Hansen and T. J. Sargent. Robust control and model uncertainty. American Economic Review, 91:60-66, 2001.
[28] G. Iyengar. Robust dynamic programming. To appear in Math. Oper. Res.. Available at http://www.corc.ieor.columbia.edu/reports/techreports/tr-2002-07.pdf, 2002.
[29] R. Jagannathan. Minimax procedure for a class of linear programs under uncertainty. Oper. Res., 25:173-177, 1977.
[30] M. J. Kearns and U. V. Vazirani. An introduction to computational learning theory. MIT Press, Cambridge, MA, 1997.
[31] P. M. Long. The complexity of learning according to two models of a drifting environment. Machine Learning, 37(3):337-354, 1999.
[32] A. Nemirovski. On tractable approximations of randomly perturbed convex constraints. In Proc. $42 n$ d IEEE Conf. Dec. Contr. (CDC), volume 3, pages 2419- 2422, 2003.
[33] A. Nemirovski and A. Shapiro. Scenario approximations of chance constraints. To appear in Probabilistic and randomized methods for design under uncertainty, 2004.
[34] S. T. Rachev. Probability metrics and the stability of stochastic models. John Wiley \& Sons, 1991.
[35] A. Ruszczynski and A. Shapiro, editors. Stochastic Programming. Handbook in Operations Research and Management Science. Elsevier, 2003.
[36] A. Ruszczynski and A. Shapiro. Optimization of risk measures. Available at http://www.optimization-online.org/DB_HTML/2004/02/822.html, 2004.
[37] A. Ruszczynski and A. Shapiro. Optimization of risk measures. Available at http://ideas.repec.org/p/wpa/wuwpri/0407002.html, 2004.
[38] A. Shapiro. Some recent developments in stochastic programming. ORB Newsletter, 13, March 2004. Available at http://www.ballarat.edu.au/ard/itms/CIAO/ORBNewsletter/issue13.shtml\#11.
[39] A. Shapiro and S. Ahmed. On a class of minimax stochastic programs. To appear in SIAM J. Opt., 2004.
[40] A. Shapiro and A. J. Kleywegt. Minimax analysis of stochastic problems. Optimization Methods and Software.
[41] V. Strassen. The existence of probability measures with given marginals. Ann. Math. Stat., 36:423-439, 1965.
[42] H. Thorisson. Coupling, Stationary, and Regeneration. Probability and its Applications. Springer-Verlag, 2000.
[43] V. N. Vapnik. The Nature of Statistical Learning Theory. Springer, New York, NY, 1995.
[44] J. Žáčková. On minimax solutions of stochastic linear programs. Čas. Pěst. Mat., pages 423-430, 1966.


[^0]:    *Submitted to Math. Prog.
    ${ }^{\dagger}$ IEOR Department, Columbia University, New York, New York 10027. Email: ee168@columbia.edu. Research partially supported by NSF grant CCR-00-09972.
    ${ }^{\ddagger}$ IEOR Department, Columbia University, New York, New York 10027. Email: garud@columbia.edu. Research partially supported by NSF grants CCR-00-09972, DMS-01-04282 and ONR grant N000140310514.

