Randomized Sampling for Large Zero-Sum Games[☆,☆☆]

Shaunak D. Bopardikar[a,*], Alessandro Borri[b], João P. Hespanha[c], Maria Prandini[d], Maria D. Di Benedetto[e]

[a] *United Technologies Research Center Inc., Berkeley, CA, USA*
[b] *Istituto di Analisi dei Sistemi ed Informatica "A. Ruberti", Consiglio Nazionale delle Ricerche (IASI-CNR), Rome, Italy.*
[c] *Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, CA 93106, USA*
[d] *Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy*
[e] *Department of Electrical and Information Engineering, University of L'Aquila, Italy*

Abstract

This paper addresses the solution of large zero-sum matrix games using randomized methods. We formalize a procedure, termed as the *sampled security policy (SSP) algorithm*, by which a player can compute policies that, with a high confidence, are security policies against an adversary using randomized methods to explore the possible outcomes of the game. The SSP algorithm essentially consists of solving a stochastically sampled subgame that is much smaller than the original game. We also propose a randomized algorithm, termed as the *sampled security value (SSV) algorithm*, which computes a high-confidence security-level (i.e., worst-case outcome) for a given policy, which may or may not have been obtained using the SSP algorithm. For both the SSP and the SSV algorithms we provide results to determine how many samples are needed to guarantee a desired level of confidence. We start by providing results when the two players sample policies with the same distribution and subsequently extend these results to the case of mismatched distributions. We demonstrate the usefulness of these results in a hide-and-seek game that exhibits exponential complexity.

Keywords: Game theory, Randomized algorithms, Zero-Sum Games, Optimization

1. Introduction

This paper addresses zero-sum games in which one or both players are faced with a large number of choices, possibly infinitely many. For such games, the computation of security levels (i.e., worst-case outcomes) and the corresponding security policies requires the exploration of very large decision trees.

Games where players are faced with deciding among a very large number of options arise in combinatorial problems, where the number of possible options grows exponentially with the size of the problem. This situation is common to many domains: In *path planning*, the number of possible decisions typically increases exponentially with the number of points to visit [cf., e.g., Bellman (1962)]. In *network security*, system administrators need to consider multi-stage, multi-host attacks that may consist of long sequences of actions by an attacker in their attempt to circumvent the system defenses [cf., e.g., Lye and Wing (2005)]. In practice, this leads to policy spaces that grow exponentially with the number of stages involved in an attack. More generally, in *partial-information feedback games* players must choose feedback policies that assign an action to each possible observation and therefore the number of feedback policies grows exponentially with the size of the players' observation spaces [cf., e.g., Hespanha and Prandini (2001); Bopardikar and Hespanha (2011)].

The exploration of large policy spaces is generally a hard task that can become computationally intractable when addressing partial information games [NP-complete in the size of the game tree, see Frank and Basin (1999)]. A method to face this issue is Monte Carlo sampling. The key idea is to confine the search to a decision tree of reduced size by guessing or sampling the other player's moves, and then use conventional minimax search to determine the strategy to play with. Techniques based on this idea has been successfully applied to several partial information games such as, e.g., Scrabble [cf. Frank (1989)], Bridge [cf. Ginsberg (1996)], and Kriegspiel chess [cf. Parker et al. (2005)]. The survey paper Browne et al. (2012) shows how Monte Carlo sampling has becoming increasing popular and has been extensively adopted, not only in the game context, but also in other domains such as e.g. path planning. Indeed, when the underlying system is stochastic

[☆] This material is based upon work supported in part by ARO MURI Grant number W911NF0910553, in part by the Center of Excellence for Research DEWS, University of L'Aquila, Italy, and in part by the European Commission under the MoVeS project, FP7-ICT-2009-257005.

[☆☆] A preliminary version of this work entitled "Randomized Sampling for Large Zero-Sum Games" was presented at the 2010 IEEE Conference on Decision and Control, Atlanta, GA, USA.

[*] Corresponding author. This work was performed while this author was at the University of California Santa Barbara, CA, USA.

Email addresses: bshaunak@gmail.com (Shaunak D. Bopardikar), alessandro.borri@iasi.cnr.it (Alessandro Borri), hespanha@ece.ucsb.edu (João P. Hespanha), prandini@elet.polimi.it (Maria Prandini), mariadomenica.dibenedetto@univaq.it (Maria D. Di Benedetto)

but it is difficult to derive an analytic description of the probabilistic distribution characterizing its evolution, solutions based on simulation are typically adopted, which entails the use of Monte Carlo sampling.

The recent successes in using randomized methods to explore large decision trees [e.g., in LaValle and Kuffner (2001); Hinton et al. (2006); Browne et al. (2012)] motivates the question that is behind the results in this paper: *Suppose that my opponent is using a randomized algorithm to explore the game decision tree, can I produce a security value and an associated security policy that are correct with high probability?* The answer to this question is affirmative and we show that such security values/policies can be constructed using randomized algorithms. What is somewhat surprising about the results reported is that one can obtain high-confidence security policies by restricting ones attention to a subset of policies that can be much smaller than the total set of policies available to the players. Moreover, this restricted set of policies may be quite different from the set of policies that the opponent considered in her randomized exploration of the game decision tree.

We call *sampled security policy* (SSP) the randomized algorithm proposed to obtain high-confidence security values/policies. The SSP algorithm can be described as follows: Suppose that player P_2 selected a policy based on a random exploration of the policies available to both players. The precise algorithm used by P_2 to select her policy based on this extraction is typically unknown. Player P_1 should then proceed as follows: randomly select a subset of the total set of available policies to both players; construct the zero-sum matrix game corresponding to the selected subset of policies, *ignoring all other policies available to the two players*; and compute the security value/policies associated with the matrix. Player P_1 can select either a mixed or a pure value/policy. In both cases, since a large number of policies have been ignored, the security policies obtained by this process will generally not be security policies for the whole game and therefore player P_1 may obtain an outcome that is strictly worse than the value computed based on her submatrix. However, we show that this happens with low probability as long as the size of the submatrix is sufficiently large. Moreover, this result holds regardless of the algorithm used by P_2 to compute her policy based on the random tree exploration. In fact, P_2 could also be using the SSP algorithm to compute her own policy.

Related Work

Two-player zero-sum matrix games have been studied extensively over the past decades [cf. the textbook by Basar and Olsder (1999)]. The classical Mini-Max theorem [cf. Von Neumann (1928)] guarantees the existence of an optimal pair of strategies for the two players, each of which is a security policy for the corresponding player. However, when the matrix is of large size, the computation of the optimal strategies involves solving optimization problems with a large number of variables and constraints.

A probabilistic approach has proven to be computationally efficient in evaluating games with large sizes. Using probabilistic analysis, the existence of simple, near-optimal strategies over a subset with logarithmically smaller size of the original matrix game was established in Lipton and Young (1994). A popular method to solve win-lose type of multi-stage or dynamic games is to evaluate the root of a game tree, in which every node is alternately an *AND* and an *OR* operation, while the leaves have a value of either 0 or 1. Motwani and Raghavan (1995) present randomized algorithms to evaluate such game trees more efficiently than using deterministic algorithms.

Randomized methods have been successful in providing efficient solutions to complex control design problems with probabilistic guarantees. Khargonekar and Tikku (1996) adopt a probabilistic approach to show the existence of randomized algorithms with polynomial complexity to solve complex robust stability analysis problems. Tempo et al. (1997) propose a randomized method for a probabilistic analysis of the worst-case controller performance, and determine sample size bounds. More recently, Tempo et al. (2004) discuss the application of randomized methods to several control design problems in the presence of uncertainty. Randomized methods have also been used to provide a probabilistic approximation to the minimax value of a cost function in robust control design problems [cf. Fujisaki and Kozawa (2003)]. Their sample complexity requirement is, in general, much higher than for the notion of security that we propose in this paper, since they are concerned with the sampled minimax value being close to the global minimax value with high confidence. A randomized approach is used in the linear programming reformulation of approximate dynamic programming in de Farias and Roy (2004). Vidyasagar (1998); Vidyasagar and Blondel (2001) demonstrate the use of randomized algorithms to solve control design problems and a number of well known complex problems in matrix theory through a statistical learning approach. Statistical learning theory [cf. Vapnik (1998)] provides a framework for probabilistic robust control synthesis. Using these tools, Alamo et al. (2009) consider semi-infinite optimization problems under uncertainty with a possibly non-convex objective function.

In Calafiore and Campi (2006); Campi and Garatti (2008); Campi and Calafiore (2009), the authors introduce the so-called *scenario approach* to solve convex optimization problems with an infinite number of constraints. Possible applications of this approach to systems and control are discussed in Calafiore and Campi (2006) and in Campi et al. (2009). Calafiore (2009) and Calafiore (2010) study the sample complexity of randomized approaches to system analysis and design, and provide, in particular, an explicit expression of the sample-size for the scenario approach to convex optimization based on an approximation of the implicit expression given in Campi and Garatti (2008). These bounds were further refined in Alamo et al. (2010). The results in these papers are instrumental to

establish several of the results in the present paper.

Contributions

Throughout the paper, we explain the results from the perspective of the player P_1 — the minimizer, — who finds herself playing against an opponent P_2 — the maximizer, — who computes her policy based on a random exploration of the game decision tree.

The contributions of this paper are four-fold. First, we show that the SSP algorithm provides a security policy for P_1 with probability $1-\delta$, provided that the size of the subgame solved by P_1 is sufficiently large. We provide two bounds on the size of the subgame, one that is valid when P_1 uses general mixed policies and the other when P_1 is restricted to consider only pure policies. The latter may require much smaller submatrix sizes when the entries of the matrix A take values in a finite set. The bounds are *game independent* and can be computable a-priori for any desired confidence level $1-\delta$, $\delta > 0$. While the size of the subgame grows with the desired confidence level $1-\delta$, it is *completely independent of the size of the original matrix game*, which could, in fact, be even infinite and not even have a value. Moreover, this bound is also independent of the precise algorithm that P_2 uses to construct her policy based on the portion of the tree that she explored.

The results outlined above assume that, while P_1 does not know the precise subtree that P_2 explored to compute her policy, P_1 does know the distribution that P_2 used to construct her random subtree. When this is not the case, there will be a mismatch between the distribution that P_1 uses in the SSP algorithm and the distribution that P_2 uses for her random exploration. The second contribution of the paper addresses this issue via two approaches. The first approach adopts sample complexity bounds obtained in Erdoğan and Iyengar (2006), which deals with the so-called ambiguous chance constrained problems. More precisely, we determine bounds on the sizes of the submatrices in the SSP algorithm when the mismatch between the distributions used by the two players remains below a specified distance $\rho < 1$, measured in the Prohorov metric. This approach requires no knowledge of the matrix game, but the bounds hold only when the confidence parameter satisfies the condition $\rho < \delta$, for a desired confidence level of $1-\delta$. The second approach is based on a novel characterization of the distance between the sampling distributions, which we call the *mismatch factor*, and is applicable to any confidence level $1-\delta$ and any mismatch factor between the distributions. However, as one would expect, for a given confidence lever, a large mismatch factor requires a large number of samples. These results take advantage of the game structure and, in fact, when the mismatch is associated with policy domination, we show that the bounds with mismatch are exactly the same as the ones without mismatch. Essentially, this result states that if P_1 knows that a particular subset $S_{\text{worse-for-}\mathsf{P}_2}$ of P_2's policies is dominated by another subset of P_2's policies $S_{\text{better-for-}\mathsf{P}_2}$ (in the sense that $S_{\text{worse-for-}\mathsf{P}_2}$ is worse than $S_{\text{better-for-}\mathsf{P}_2}$ from P_2's perspective), then P_1 need not sample policies from $S_{\text{worse-for-}\mathsf{P}_2}$. The usefulness of the second approach goes beyond investigating the confidence of the SSP algorithm, as it extends the bounds of the scenario approach derived in Campi and Calafiore (2009) and Alamo et al. (2010) to mismatched distributions.

Third, we propose a randomized algorithm, which we call *sampled security-value* (SSV) that P_1 can use to obtain a high-probability security level for a given policy. The bound on the size of the subgame that P_1 extracts to determine her high-probability security level holds for any policy available to P_1, regardless of whether or not this policy was obtained from the SSP algorithm. As for the SSP algorithm, the computation required by the SSV algorithm is independent of the size of the original matrix game and also of the precise algorithm that P_2 uses to construct her policy. When applied to a policy obtained using the SSP algorithm for a confidence level δ_{SSP}, the SSV algorithm can be used to study the security of the policy for different (perhaps tighter) confidence levels δ.

Fourth and finally, we apply the SSP and SSV algorithms to solve a hide-and-seek game, in which one player hides a treasure in one of N points and the other player searches for the treasure by visiting each of the points. This is formalized as a zero-sum game in which the player that hides the treasure wants to maximize the distance that the other player needs to travel until the treasure is found. To determine the optimal strategy for this game, one would need to solve a matrix game whose size is $N \times N!$. Thus, exact solutions to this problem require computation that scales exponentially with the number of points N. Our approach is *independent of the size of the game* and therefore the size of the matrix plays no role in the amount of computation required.

As compared to the preliminary conference version [cf. Bopardikar et al. (2010)], this paper presents new results that include the version of the SSP algorithm for pure policies and its analysis, the mismatch in the distributions used by the players to construct the subgames, and improves upon the *explicit* sample size bounds using the results in Alamo et al. (2010). The more recent paper [cf. Bopardikar and Hespanha (2011)] formalizes the sampling procedure to dynamic or multi-stage, partial information games.

Organization

This paper is organized as follows. The problem formulation and the SSP algorithm are presented in Section 2. Bounds on the size of the subgame to provide high-confidence SSP solutions are established in Section 3 for the case when the two players use identical distributions to sample the matrix. These bounds are extended in Section 4 to allow mismatch between the distributions. Section 5 presents the SSV algorithm and the related bounds on the size of the subgame. Finally, we demonstrate the

procedure applied to the hide-and-seek problem in Section 6.

2. Problem Formulation

Consider a zero-sum matrix game defined by an arbitrary $M \times N$ real-valued matrix A, in which player P_1 is the minimizer and selects rows and player P_2 is the maximizer and selects columns. We are interested in problems where the number N of (pure) policies available to P_2 is very large, typically due to combinatorial explosion, forcing P_2 to explore a random subset of her own (pure) policy space with only $n_2 \ll N$ policies, and perhaps also only a random subset of the possibles responses by P_1. Based on this, P_2 selects a policy z^* that she will use to play against P_1.

Denoting by $\mathcal{B}^{k\times\ell}$ the set of $k \times \ell$ left-stochastic $(0,1)$-valued matrices (i.e., matrices whose entries belong to the set $\{0,1\}$ with exactly one 1 per column), we can express the process by which P_2 samples her own policy space by selecting a random matrix Π_2 from the set $\mathcal{B}^{N\times n_2}$. The matrix $\Pi_2 \in \mathcal{B}^{N\times n_2}$ has one row for each of the possible policies of P_2 in the original game defined by A and one column for each policy that was actually explored by P_2. A one in row i, column j of Π_2 signifies that the jth policy explored by P_2 corresponds to the column i of A. P_2's random exploration results in a mixed policy z^* that can be written as

$$z^* = \Pi_2 z_2^* \in \mathcal{S}_N, \quad z_2^* \in \mathcal{S}_{n_2}, \tag{1}$$

where, for a given integer k, $\mathcal{S}_k$ denotes the probability simplex of size k. P_1 may know the distribution used to extract Π_2, but *will not know the matrix* Π_2 *that was actually extracted nor which algorithm was used to determine* z_2^* and therefore will not know the policy z^* obtained by P_2.

For P_1 to compute a high-confidence response against P_2's policy z^* in (1), we introduce the *sampled security policy (SSP) Algorithm* 1.

For the SSP algorithm to be useful, it needs to provide appropriate guarantees of correctness, which are formalized by the following definitions. We say that *the SSP algorithm is* ϵ*-secure for player* P_1 *with confidence* $1-\delta$ if

$$\mathrm{P}_{\Gamma_1,\Pi_1,\Pi_2}\left(y^{*\prime} A z^* \le \bar{V}(A_1) + \epsilon\right) \ge 1-\delta. \tag{3}$$

Here and in the sequel, we use a subscript in the probability measure P to remind the reader which random variables define the event that is being measured. In essence, condition (3) states that the probability that the outcome of the game will violate P_1's sampled security value $\bar{V}(A_1)$ by more than ϵ is smaller than δ. As stated, this definition requires the bound to hold regardless of the algorithm used by P_2 to select her policy z^*. In fact, we even allow z^* to be obtained using an algorithm that randomly explores

Algorithm 1 [SSP Algorithm]

1: P_1 randomly selects m_1 rows and n_1 columns of A, which she uses to construct an $m_1 \times n_1$ submatrix A_1 of A. This can expressed by the selection of two random matrices $\Gamma_1 \in \mathcal{B}^{M\times m_1}$ and $\Pi_1 \in \mathcal{B}^{N\times n_1}$ and then computing the product $A_1 = \Gamma_1' A \Pi_1$.

2: P_1 computes the mixed security value $\bar{V}(A_1)$ and the corresponding security policy y_1^* for A_1:

$$\bar{V}(A_1) = \max_{z\in\mathcal{S}_{n_1}} y_1^{*\prime} A_1 z = \min_{y\in\mathcal{S}_{m_1}} \max_{z\in\mathcal{S}_{n_1}} y' A_1 z \tag{2}$$

We call $\bar{V}(A_1)$ P_1*'s sampled security value.* When multiple security policies y_1^* exist, P_1 selects for y_1^* the one with the minimum Euclidean norm (since the set of security policies is convex, it contains a unique element with minimum norm).

3: P_1 computes her mixed policy for the original game:

$$y^* := \Gamma_1 y_1^*,$$

resulting in the outcome $y^{*\prime} A z^* = y_1^{*\prime} \Gamma_1' A \Pi_2 z_2^*$. We call y^* P_1*'s sampled security policy.*

P_1's policy space[1]. While our results do not depend on it, P_2 could have obtained z^* also using the SSP algorithm.

The previous definition guarantees that P_1 will be surprised with (low) probability δ when playing with policies obtained from a one-shot solution to the SSP algorithm. However, no specific guarantee is given regarding the inherent safety of the specific policy y^* obtained using the SSP algorithm. So, e.g., suppose that player P_1 computes y^* once using the SSP algorithm and then plays this policy multiple times against a sequence of policies z^* for P_2, each obtained by a distinct random explorations of her policy space. Then P_1 could conceivably be surprised many more times than one would expect for a low value of δ. This would happen if she was "unlucky" and got a (low probability) y^* that is particularly bad or a value $\bar{V}(A_1)$ that is particularly optimistic. To avoid this scenario, we introduce an additional notion of security that refers to the security of a specific policy/value: we say that *a policy* y^* *with value* $\bar{V}(A_1)$ *is* ϵ*-secure for player* P_1 *with confidence* $1-\delta$ if

$$\mathrm{P}_{\Pi_2}\left(y^{*\prime} A z^* \le \bar{V}(A_1) + \epsilon \mid y^*,\ \bar{V}(A_1)\right) \ge 1-\delta. \tag{4}$$

Note that the subscript in the probability measure now only includes the matrix corresponding to randomized exploration of the policy space by P_2 since the probability guarantees are given for a specific security policy and value of P_1.

So far, we have not specified the joint distribution of the row/column extraction matrices Γ_1 and Π_1 for P_1 in

[1] In this case, the probability measure in (3) depends on additional random variables that we do not explicitly include in the subscript.

the SSP Algorithm 1, but these distributions, jointly with that of matrix Π_2 for P_2, clearly affect the outcome of the algorithm. In the context of noncooperative games, one should presume the extractions of the two players to be independent of each other. For simplicity, we further assume that players extract rows and columns independently, as stated in the following assumption:

Assumption 2.1 (Independence) *The random matrices Γ_1 and Π_1 in the SSP Algorithm 1 and the matrix Π_2 corresponding to player P_2's randomized exploration are statistically independent and each of them has independent and identically distributed columns.* □

Under Assumption 2.1, we shall determine in Section 3 bounds on the size of the random matrices extracted by P_1 in the SSP Algorithm 1 that guarantee high-confidence ϵ-security results with $\epsilon = 0$. These results are valid when P_1 knows precisely the distribution used by P_2 to explore her game decision tree, i.e., to extract columns of the game matrix A. If we allow for mismatched distributions, we can then prove ϵ-security results with a value for $\epsilon > 0$ that depends on the distance between the distributions used by P_1 and P_2 to extract columns of A (see Section 4).

Remark 2.1 (General games) The results in this paper do not depend on the fact that the original game is a finite matrix game. They extend trivially to any cost-function $J(u, d)$, $u \in \mathcal{U}$, $d \in \mathcal{D}$ where $\mathcal{U}$ and $\mathcal{D}$ denote the sets of policies for the minimizer and maximizer, respectively. In fact, it is not even necessary that the original game has saddle-point policies since all that the results use is the fact that, when we take finite samples of the sets of policies, we obtain finite matrix games. In fact, these results also apply to dynamic games, as is discussed in Bopardikar and Hespanha (2011). □

3. Bounds for Probabilistic Guarantees

In this section, we present theoretical bounds on the number of policies that player P_1 needs to consider for the SSP Algorithm to guarantee desired confidence levels. The results in this section refer to the case where the players sample policies for P_2 (i.e., *columns* of A) using identical distributions. This assumption is subsequently relaxed in Section 4.

3.1. Mixed Sampled Security Policies

The main result of this section provides a bound on the size of the submatrix A_1 in the SSP Algorithm that guarantees ϵ-security with $\epsilon = 0$ for the mixed policy y^*. We recall that P_2 is assumed to use a policy z^* of the form (1), where Π_2 is a column-selection matrix and z_2^* some vector in $\mathcal{S}_{n_2}$ that is obtained using a deterministic or stochastic algorithm. The case of P_2 using a sample of P_1's policies to determine her policy z^* also gets included as none of the results in this paper require P_2 to use the same distribution as that used by P_1 for extracting rows of A.

Theorem 3.1 (SSP Algorithm) *Suppose that Assumption 2.1 holds and that $\Pi_1 \in \mathcal{B}^{N\times n_1}$ and $\Pi_2 \in \mathcal{B}^{N\times n_2}$ have identically distributed columns. The SSP Algorithm is $(\epsilon = 0)$-secure for P_1 with confidence $1-\delta$, $\delta \in (0,1)$ as long as*[2]

$$n_1 = \left\lceil \frac{m_1 + 1}{\delta} - 1 \right\rceil \bar{n}_2, \tag{5}$$

with $\bar{n}_2 \geq n_2$. Additionally, suppose that we increase n_1 to satisfy

$$n_1 = \left\lceil \frac{1}{\delta}\left(m_1 + \sqrt{2m_1 \ln \frac{1}{\beta}} + \ln \frac{1}{\beta}\right) \right\rceil \bar{n}_2, \tag{6}$$

for some $\beta \in (0,1)$. Then, with probability larger than $1-\beta$, the SSP Algorithm generates a sampled security policy y^* with value $\bar{V}(A_1)$ that is $(\epsilon = 0)$-secure for P_1 with confidence $1-\delta$. □

In words, this result states that it is always possible to guarantee $(\epsilon = 0)$-security for P_1, if she constructs her submatrix A_1 utilizing a sufficiently large number of columns n_1. In particular, she always needs to choose a number of columns n_1 larger than the number of columns n_2 that P_2 is considering for her mixed policies [cf. (5) and (6)]. The additional number of columns that P_1 needs to consider is a function of the number m_1 of rows that P_1 wants to consider for her mixed policy and the desired confidence level.

The probability $1-\beta$ associated with y^*'s security probabilistic guarantee accounts for the possibility that the confidence bound (4) fails altogether due to an "unfortunate" sample used by P_1 to compute y^*. However, note that only the logarithm of the confidence level β appears (6) and therefore a relatively small value for the number of columns n_1 suffices to make β extremely small (and, hence, $1-\beta \simeq 1$).

Remark 3.2 (P_1's knowledge of n_2) According to Theorem 3.1, for player P_1 to enjoy guaranteed $(\epsilon = 0)$-security with confidence $1-\delta$, she must know an upper bound $\bar{n}_2$ on the number of columns that P_2 used to construct the policy z^* in (1). However, if P_1 does not know $\bar{n}_2$ precisely and, e.g., underestimates $\bar{n}_2$ by a certain percentage, then (5) and (6) are still useful in the sense that they predict that the performance degradation in the actual confidence level δ grows proportionately with $\bar{n}_2$. This is because the bounds in (5) and (6) essentially scale with $\frac{\bar{n}_2}{\delta}$. □

[2] Given a scalar $x \in \mathbb{R}$, we denote by $\lceil x \rceil$ the smallest integer that is larger than or equal to x.

Proof of Theorem 3.1: By definition of the security value $\bar{V}(A_1)$, we have that

$$\begin{aligned}\bar{V}(A_1) &= \min_{y\in\mathcal{S}_{m_1}} \max_{z\in\mathcal{S}_{n_1}} y'\Gamma_1'A\Pi_1 z \\ &= \min_{y\in\mathcal{S}_{m_1}} \max_{j\in\{1,\ldots,n_1\}} y'\Gamma_1'A\Pi_1 e_j(n_1) \\ &= \min_{\theta\in\Theta}\Big\{v : y'\Gamma_1'A\Pi_1 e_j(n_1)\le v,\ \forall j\in\{1,\ldots,n_1\}\Big\}, \quad (7)\end{aligned}$$

where $\theta := (y_1, v)$, $\Theta := \mathcal{S}_{m_1}\times\mathbb{R}$, and we use $e_j(n)$ to denote the jth element of the canonical basis of $\mathbb{R}^n$.

Since n_1 is an integer multiple of $\bar{n}_2$, i.e., $n_1 = K\bar{n}_2$ with $K = \left\lceil \frac{m_1+1}{\delta} - 1\right\rceil$, we can use the $K\bar{n}_2$ columns of $\Pi_1 \in \mathcal{B}^{N\times K\bar{n}_2}$ to construct K independent and identically distributed (i.i.d.) matrices $\Delta_1, \Delta_2, \ldots, \Delta_K \in \mathcal{B}^{N\times\bar{n}_2}$. For an arbitrary realization of the matrix $\Gamma_1 \in \mathcal{B}^{M\times m_1}$, which is independent of the Δ_i by Assumption 2.1, let us define the function $f_{\Gamma_1} : \Theta\times\mathcal{B}^{N\times\bar{n}_2}\to\mathbb{R}$ by

$$f_{\Gamma_1}(\theta,\Delta) = \max_{j\in\{1,\ldots,\bar{n}_2\}} y_1'\Gamma_1'A\Delta e_j(\bar{n}_2) - v. \quad (8)$$

We can then rewrite (7) as

$$\bar{V}(A_1) = \min_{\theta\in\Theta}\big\{v : f_{\Gamma_1}(\theta,\Delta_i)\le 0,\ \forall i\in\{1,\ldots,K\}\big\} \quad (9)$$

and conclude from (Campi and Calafiore, 2009, Proposition 3) that the (conditional) probability that another matrix Δ sampled independently from the same distribution as the Δ_i satisfies the constraint $f_{\Gamma_1}(\theta^*,\Delta)\le 0$ can be lower-bounded as follows:

$$\mathrm{P}_{\Pi_1,\Delta}\left(f_{\Gamma_1}(\theta^*,\Delta)\le 0 \mid \Gamma_1\right) \ge \frac{K-m_1}{K+1} \ge 1-\delta, \quad (10)$$

where θ^* denotes the value in Θ that achieves the minimum in (9) and the second inequality is a consequence of (5). Since the minimum in (9) is achieved for the sampled security policy/value $\theta^* = (y_1^*, \bar{V}(A_1))$, we can use the definition (8) of f_{Γ_1} to re-write (10) as

$$\mathrm{P}_{\Pi_1,\Delta}\left(y_1^{*\prime}\Gamma_1'A\Delta e_j(\bar{n}_2)\le \bar{V}(A_1), \quad \forall j\in\{1,\ldots,\bar{n}_2\}\mid\Gamma_1\right)\ge 1-\delta.$$

Since $n_2\le\bar{n}_2$, we further conclude that

$$\mathrm{P}_{\Pi_1,\Delta}\left(y_1^{*\prime}\Gamma_1'A\Delta e_j(n_2)\le \bar{V}(A_1), \quad \forall j\in\{1,\ldots,n_2\}\mid\Gamma_1\right)\ge 1-\delta.$$

Under Assumption 2.1, when the columns of Π_1 and Π_2 are identically distributed, the matrix consisting of the first n_2 columns of Δ can be viewed as the matrix Π_2 and the inequality above implies that

$$\mathrm{P}_{\Pi_1,\Pi_2}\left(y_1^{*\prime}\Gamma_1'A\Pi_2 e_j(n_2)\le \bar{V}(A_1), \quad \forall j\in\{1,\ldots,n_2\}\mid\Gamma_1\right)\ge 1-\delta.$$

Since

$$\begin{aligned}y_1^{*\prime}\Gamma_1'A\Pi_2 e_j(n_2)\le\bar{V}(A_1),\ \forall j\in\{1,\ldots,n_2\} \quad &\Rightarrow \\ \Rightarrow \quad y_1^{*\prime}\Gamma_1'A\Pi_2 z\le\bar{V}(A_1),\ \forall z\in\mathcal{S}^{n_2},\end{aligned}$$

we conclude that

$$\mathrm{P}_{\Pi_1,\Pi_2}\left(y_1^{*\prime}\Gamma_1'A\Pi_2 z_2^*\le\bar{V}(A_1)\mid\Gamma_1\right)\ge 1-\delta.$$

We have shown that this bound holds for an arbitrary realization of Γ_1, therefore it also holds for the unconditional probability, which shows that the SSP Algorithm is $(\epsilon=0)$-secure for P_1 with confidence $1-\delta$ according to (3).

If instead of using (Campi and Calafiore, 2009, Proposition 3) and (5) to obtain (10), we use (Alamo et al., 2010, Theorem 4) and (6), we obtain instead that

$$\mathrm{P}_{\Delta}\left(f_{\Gamma_1}(\theta^*,\Delta)\le 0 \mid \Gamma_1,\theta^*\right)\ge 1-\delta, \quad (11)$$

with probability larger than $1-\beta$, where the confidence level $1-\beta$ refers to the extraction of $\Pi_1 = [\Delta_1,\ldots,\Delta_K]$ that defines θ^* (given Γ_1). The proof can now proceed exactly as before, but with (10) replaced by (11), which now involves a probability conditioned to $\theta^* = (y^*, \bar{V}(A_1))$. This shows that if n_1 satisfies (6), then with probability larger than $1-\beta$, the policy y^* with value $\bar{V}(A_1)$ is $(\epsilon=0)$-secure for P_1 with confidence $1-\delta$. ■

3.2. Pure Sampled Security Policy

Suppose that P_1 restricts herself to use pure policies in Step 2 of the SSP Algorithm 1. If we let $e_i(m_1)$ denote the ith element of the canonical basis of $\mathbb{R}^{m_1}$, then Step 2 becomes:

2: P_1 computes the pure security value $\bar{V}_{\mathrm{pure}}(A_1)$:

$$\begin{aligned}\bar{V}_{\mathrm{pure}}(A_1) &= \max_{z\in\mathcal{S}_{n_1}} e_{i^*}'(m_1)A_1 z \\ &= \min_{i\in\{1,\ldots,m_1\}} \max_{z\in\mathcal{S}_{n_1}} e_i'(m_1)A_1 z,\end{aligned}$$

and the corresponding pure security policy y_1^* for A_1:

$$y_1^* = e_{i^*}(m_1).$$

We call $\bar{V}_{\mathrm{pure}}(A_1)$ P_1*'s pure sampled security value.* When multiple pure security policies y_1^* exist, P_1 can pick any of them.

A bound similar to (6) can be established for the resulting *pure SSP Algorithm*, but the number of columns n_1 that P_1 needs to sample can often be much smaller. Note that the bound still holds for any policy z^* for player P_2 of the form (1), pure or mixed.

Theorem 3.3 (Pure SSP Algorithm) *Suppose that Assumption 2.1 holds and that* $\Pi_1\in\mathcal{B}^{N\times n_1}$ *and* $\Pi_2\in\mathcal{B}^{N\times n_2}$

have identically distributed columns. Suppose that we select

$$n_1 = \left\lceil \frac{1}{\delta}\Big(\ln\big(m_1 \cdot \#(\Gamma_1' A)\big) + \ln\frac{1}{\beta}\Big)\right\rceil \bar{n}_2, \qquad (12)$$

for $\beta, \delta \in (0,1)$ and $\bar{n}_2 \geq n_2$, where $\#(\Gamma_1' A)$ denotes the total number of distinct values that the entries of $\Gamma_1' A$ can take. Then, with probability $1-\beta$, the pure SSP Algorithm generates a pure sampled security policy y^* with value $\bar{V}_{\rm pure}(A_1)$ that is $(\epsilon = 0)$-secure for P_1 with confidence $1-\delta$. □

In several matrix games the number of distinct values that entries of A can take is small and therefore $\#(\Gamma_1' A)$ is small. This occurs, e.g., in win-lose-tie games. For such games, (12) provides significant computational gains with respect to (6) because the bound in (12) grows with the *logarithm* of m_1, whereas the one in (6) grows *linearly* with m_1. Even if the matrix A can have many distinct values, significant computational savings are possible since $\#(\Gamma_1' A) \leq m_1 N$ and hence, at worse, (12) still only grows with the *logarithm* of $m_1^2 N$. The price to pay for these computational savings is that the pure security level $\bar{V}_{\rm pure}(A_1)$ could be higher than the value $\bar{V}(A_1)$ obtained with mixed policies.

The proof of Theorem 3.3 is conceptually similar to the second part of the proof of Theorem 3.1, with the main difference being that the policy selection involves optimizing over a finite set of cardinality $\#(\Gamma_1' A)$, and, hence, we can use the bounds in (Alamo et al., 2010, Theorem 3) instead of those in (Alamo et al., 2010, Theorem 4).

Proof of Theorem 3.3: Let $e_k(n)$ denote the kth element of the canonical basis of $\mathbb{R}^n$. By the definition of the pure security value $\bar{V}_{\rm pure}(A_1)$, we have that

$$\begin{aligned}
\bar{V}_{\rm pure}(A_1) &= \min_{i\in\{1,\ldots,m_1\}} \max_{z\in\mathcal{S}_{n_1}} e_i'(m_1)\Gamma_1' A \Pi_1 z \\
&= \min_{i\in\{1,\ldots,m_1\}} \max_{j\in\{1,\ldots,n_1\}} e_i'(m_1)\Gamma_1' A \Pi_1 e_j(n_1) \\
&= \min_{\theta\in\Theta} \big\{ v : e_i'(m_1)\Gamma_1' A \Pi_1 e_j(n_1) \leq v,\ \forall j \in \{1,\ldots,n_1\}\big\},
\end{aligned} \qquad (13)$$

where $\theta := (e_i(m_1), v)$,

$$\Theta := \{e_1(m_1), e_2(m_1), \ldots, e_{m_1}(m_1)\} \times \{\Gamma_1' A\},$$

and $\{\Gamma_1' A\}$ denotes the set of all distinct values that the entries of $\Gamma_1' A$ can take. The set Θ thus have $m_1 \cdot \#(\Gamma_1' A)$ distinct elements.

Since n_1 is an integer multiple of $\bar{n}_2$, i.e., $n_1 = K\bar{n}_2$ with $K = \left\lceil \frac{1}{\delta} \ln \frac{m_1 \cdot \#(\Gamma_1' A)}{\beta} \right\rceil$, we can use the $K\bar{n}_2$ columns of $\Pi_1 \in \mathcal{B}^{N \times K\bar{n}_2}$ to construct K i.i.d. matrices $\Delta_1, \Delta_2, \ldots, \Delta_K \in \mathcal{B}^{N\times \bar{n}_2}$. For an arbitrary realization of the matrix $\Gamma_1 \in \mathcal{B}^{M\times m_1}$, which is independent of the Δ_i by Assumption 2.1, let us define the function $g_{\Gamma_1} : \Theta \times \mathcal{B}^{N\times\bar{n}_2} \to \mathbb{R}$ by

$$g_{\Gamma_1}(\theta, \Delta) = \begin{cases} 0, & \max\limits_{j\in\{1,\ldots,\bar{n}_2\}} e_i'(m_1)\Gamma_1' A \Delta e_j(\bar{n}_2) - v \leq 0, \\ 1, & \text{otherwise,} \end{cases}$$

We can then rewrite (13) as

$$\bar{V}_{\rm pure}(A_1) = \min_{\theta\in\Theta} \big\{ v : \sum_{k=1}^{K} g_{\Gamma_1}(\theta, \Delta_i) \leq 0\big\}, \qquad (14)$$

and conclude from (Alamo et al., 2010, Theorem 3) that the (conditional) probability that another matrix Δ sampled independently from the same distribution as the Δ_i satisfies the constraint $g_{\Gamma_1}(\theta^*, \Delta) = 0$, where θ^* denotes the value in Θ that achieves the minimum in (14), can be lower-bounded by

$$\mathrm{P}_\Delta\left(g_{\Gamma_1}(\theta^*, \Delta) = 0 \mid \Gamma_1, \theta^*\right) \geq 1-\delta, \qquad (15)$$

with probability larger than $1-\beta$, where the confidence level $1-\beta$ refers to the extraction of $\Pi_1 = [\Delta_1, \ldots, \Delta_K]$ that defines θ^* (given Γ_1). Since the minimum in (14) is achieved for the sampled security policy/value $\theta^* = (y_1^*, \bar{V}_{\rm pure}(A_1))$, we can use the definition of g_{Γ_1} to re-write (15) as

$$\begin{aligned}
\mathrm{P}_\Delta\big(y_1^{*\prime}\Gamma_1' A \Delta e_j(\bar{n}_2) \leq \bar{V}_{\rm pure}(A_1),\ \forall j \in \{1,\ldots,\bar{n}_2\} \\
\mid \Gamma_1, y_1^*, \bar{V}_{\rm pure}(A_1)\big) \geq 1-\delta.
\end{aligned}$$

Since $n_2 \leq \bar{n}_2$, we further conclude that

$$\begin{aligned}
\mathrm{P}_\Delta\big(y_1^{*\prime}\Gamma_1' A \Delta e_j(n_2) \leq \bar{V}_{\rm pure}(A_1),\ \forall j \in \{1,\ldots,n_2\} \\
\mid \Gamma_1, y_1^*, \bar{V}_{\rm pure}(A_1)\big) \geq 1-\delta.
\end{aligned}$$

Under Assumption 2.1, when the columns of Π_1 and Π_2 are identically distributed, the matrix consisting of the first n_2 columns of Δ can be viewed as the matrix Π_2 and the inequality above implies that

$$\begin{aligned}
\mathrm{P}_{\Pi_2}\big(y_1^{*\prime}\Gamma_1' A \Pi_2 e_j(n_2) \leq \bar{V}_{\rm pure}(A_1),\ \forall j \in \{1,\ldots,n_2\} \\
\mid \Gamma_1, y_1^*, \bar{V}_{\rm pure}(A_1)\big) \geq 1-\delta.
\end{aligned}$$

Since

$$\begin{aligned}
y_1^{*\prime}\Gamma_1' A \Pi_2 e_j(n_2) \leq \bar{V}_{\rm pure}(A_1),\ \forall j \in \{1,\ldots,n_2\} \quad \Rightarrow \\
\Rightarrow \quad y_1^{*\prime}\Gamma_1' A \Pi_2 z \leq \bar{V}_{\rm pure}(A_1),\ \forall z \in \mathcal{S}^{n_2},
\end{aligned}$$

we conclude that

$$\mathrm{P}_{\Pi_2}\left(y_1^{*\prime}\Gamma_1' A \Pi_2 z_2^* \leq \bar{V}_{\rm pure}(A_1) \mid \Gamma_1, y_1^*, \bar{V}_{\rm pure}(A_1)\right) \geq 1-\delta.$$

This shows that, with probability larger than $1-\beta$, the sampled pure policy $y^* := \Gamma_1 y_1^*$ with with value $\bar{V}_{\rm pure}(A_1)$ is $(\epsilon = 0)$-secure for P_1 with confidence $1-\delta$. ■

Remark 3.4 The bound for the pure SSP Algorithm corresponding to (5) in Theorem 3.1 would be

$$n_1 = \left\lceil \frac{m_1 \cdot \#(\Gamma_1' A)}{\delta} - 1 \right\rceil \bar{n}_2,$$

which can be obtained by first deriving the single level of probability version of the bounds in (Alamo et al., 2010, Theorem 3), following similar steps as in (Campi and Calafiore, 2009, Proposition 3). Given that this bound is worse than (5), there is no computational advantage for player P_1 in considering pure rather than mixed policies. Consequently, this result is not included in Theorem 3.3. □

4. Mismatch in the Sampling Distributions

We now investigate the effect of a mismatch between the distribution that P_1 uses to select the columns of the policy-selection matrix Π_1 used in the SSP Algorithm of Section 3 and the distribution that P_2 uses to select the columns of the policy-selection matrix Π_2 that she uses to determine the policy z^* in (1).

We pursue two approaches: The first one is based on a characterization of the mismatch between distributions using the Prohorov metric and provides bounds that are independent of the game matrix A. The second approach provides a novel characterization of the mismatch between the two distributions that can take the matrix A into consideration. The proofs for the results in this section are deferred to the Appendix to facilitate the presentation.

4.1. Prohorov Metric-Based Approach

For a given integer k, the distributions used to select the columns of Π_1 and Π_2 can be used to construct two measures m^k and $\tilde{m}^k$, respectively, for the column-selection matrices taking values in $\mathcal{B}^{N\times k}$ with i.i.d. columns. Using the discrete metric

$$d(x_1, x_2) = \begin{cases} 1, & x_1 \neq x_2, \\ 0, & x_1 = x_2, \end{cases} \qquad \forall x_1, x_2 \in \mathcal{B}^{N\times k}, \quad (16)$$

we can regard $\mathcal{B}^{N\times k}$ as a metric space, for which the Prohorov metric between m^k and $\tilde{m}^k$ is simply given by the total variation metric

$$\pi(m^k, \tilde{m}^k) = \sup_{B\in\mathcal{F}} |m^k(B) - \tilde{m}^k(B)|, \quad (17)$$

where $\mathcal{F}$ denotes the Borel sigma-algebra on $\mathcal{B}^{N\times k}$. The following theorem is based on the results for the ambiguous chance constrained problems in Erdoğan and Iyengar (2006) and should be viewed as a generalization of the bound in (6) Theorem 3.1 for the case of mismatched distributions.

Theorem 4.1 (Matrix-independent mismatch) *Suppose that Assumption 2.1 holds, that*

$$\pi\big(m^{\bar{n}_2}, \tilde{m}^{\bar{n}_2}\big) \leq \rho < 1$$

for some $\bar{n}_2 \geq n_2$, and that we select

$$n_1 = \left\lceil 2(m_1+1) + \frac{2}{\delta-\rho}\ln\frac{1}{\beta} + \frac{2(m_1+1)}{\delta-\rho}\ln\frac{2}{\delta-\rho} \right\rceil \bar{n}_2 \quad (18)$$

for $\beta \in (0,1)$, $\delta \in (\rho, 1)$. Then, with probability larger than $1-\beta$, the SSP Algorithm generates a sampled security policy y^* with value $\bar{V}(A_1)$ that is $(\epsilon = 0)$-secure for P_1 with confidence $1-\delta$. □

In the spirit of Theorem 3.1, the bound (18) is completely independent of the matrix game A. However, this result has the limitation that it is applicable only for confidence levels $\delta > \rho$ and therefore does not permit confidence levels larger than $1-\rho$.

4.2. Mismatch Factor-Based Approach

Our second approach to characterize the impact of a mismatch between the sampling distributions relies on generalizations of the bounds of the scenario approach to convex optimization in Campi and Calafiore (2009) and Alamo et al. (2010) that were instrumental to the proof of Theorem 3.1. We start by presenting these generalizations, which are useful beyond the context of the problem considered here.

4.2.1. Scenario Optimization

Consider a sequence of K i.i.d. random variables Δ_1, Δ_2, ..., Δ_K taking values in a set D. These random variables are used to specify a set of K constraints in the following convex optimization problem:[3]

$$\theta^* = \arg\min_{\theta\in\Theta} \Big\{ c'\theta : f(\theta, \Delta_i) \leq 0, \, \forall i \in \{1, \ldots, K\} \Big\}, \quad (19)$$

where $c \in \mathbb{R}^{n_\theta}$ and the constraint-defining function $f : \Theta \times D \to \mathbb{R}$ is convex with respect to the first argument, for each fixed value of D, and Θ is a convex subset of $\mathbb{R}^{n_\theta}$.

The results that follow provide bounds on the probability that an additional independent random variable $\bar{\Delta}$, also taking values in D but with a different distribution, satisfies the following (somewhat relaxed) version of the constraint that appears in (19) for the optimal θ^*:

$$f(\theta^*, \bar{\Delta}) \leq \epsilon,$$

for some $\epsilon \geq 0$. (Campi and Calafiore, 2009, Proposition 3) and (Alamo et al., 2010, Theorem 4) provide such bounds when Δ_i and $\bar{\Delta}$ have the same distribution and $\epsilon = 0$. Denoting by m and $\tilde{m}$ the measures associated with the distributions of the Δ_i and $\bar{\Delta}$, respectively, we define the *mismatch factor between m and $\tilde{m}$* by

$$\mu_f(\epsilon) := \inf_{\mu\in\mathbb{R}} \Big\{ \mu : \tilde{m}\big(f(\theta, \bar{\Delta}) > \epsilon\big) \leq \mu \, m\big(f(\theta, \Delta) > 0\big), \quad \forall \theta \in \Theta \Big\} \quad (20)$$

[3] In case of several possible multiple minima, the one with the smallest Euclidean norm should be selected.

When $m = \tilde{m}$, we have $\mu_f(\epsilon) \leq 1$, with equality when $\epsilon = 0$. However, when the distributions do not match, $\mu_f(\epsilon)$ can be arbitrarily large. As the name indicates, the mismatch factor $\mu_f(\epsilon)$ can be viewed as a measure of how much the distributions of Δ and $\bar{\Delta}$ differ. Aside from not being a metric, it differs more fundamentally from the Prohorov metric (17) in that (i) (20) only regards the discrepancy between values of the measures for "violation" events of the type $f(\theta, \bar{\Delta}) > \epsilon$; (ii) it considers a kind of multiplicative uncertainty in the probabilities (instead of differences); and (iii) it allows for the "relaxation" parameter $\epsilon > 0$ that can bring $\mu_f(\epsilon)$ down if we are willing to allow $f(\theta, \bar{\Delta})$ to grow as large as $\epsilon > 0$. As we shall see shortly, smaller values of $\mu_f(\epsilon)$ lead to smaller probabilities of violation.

The following two results generalize (Campi and Calafiore, 2009, Proposition 3) and (Alamo et al., 2010, Theorem 4), respectively, for mismatched distributions and $\epsilon > 0$.

Lemma 4.2 *For every $\epsilon \geq 0$,*

$$\mathrm{P}_{\bar{\Delta},\Delta_1,\ldots,\Delta_K}\left(f\big(\theta^*, \bar{\Delta}\big) \leq \epsilon\right) \geq 1 - \frac{\mu_f(\epsilon)\, n_\theta}{K+1}. \tag{21}$$

□

Lemma 4.3 *Given $\epsilon > 0$, $\delta \in (0,1)$, $\beta \in (0,1)$, and*

$$K \geq \left\lceil \frac{\mu_f(\epsilon)}{\delta}\Big((n_\theta - 1) + \sqrt{2(n_\theta - 1)\ln\frac{1}{\beta}} + \ln\frac{1}{\beta}\Big)\right\rceil,$$

we have that

$$\mathrm{P}_{\bar{\Delta}}\left(f(\theta^*, \bar{\Delta}) \leq \epsilon \,|\, \theta^*\right) \geq 1 - \delta, \tag{22}$$

with probability[4] larger than $1 - \beta$. □

4.2.2. Probabilistic Guarantees

Lemmas 4.2 and 4.3 allow us to generalize Theorem 3.1 for the case of mismatched distributions. This generalization involves a family of functions $f_\Gamma : \Theta \times B^{N\times\bar{n}_2} \to \mathbb{R}$, with $\Theta := (\mathcal{S}_{m_1}, \mathbb{R})$ and $\bar{n}_2$ an integer larger than n_2, parameterized by the matrix $\Gamma \in \mathcal{B}^{M\times m_1}$ and defined by

$$f_\Gamma(\theta, \Delta) = \max_{j\in\{1,\ldots,\bar{n}_2\}} y_1'\Gamma' A \Delta e_j(\bar{n}_2) - v. \tag{23}$$

We shall use these functions to compute the mismatch factor between the measures m^k and $\tilde{m}^k$ for column-selection matrices taking values in $\mathcal{B}^{N\times k}$ with i.i.d. columns, constructed using the distributions used to select the columns of Π_1 and Π_2, respectively.

Theorem 4.4 (Matrix-dependent mismatch) *Suppose that Assumption 2.1 holds, and that*

$$\tilde{m}^{\bar{n}_2}\big(f(\theta, \bar{\Delta}) > \epsilon\big) \leq \mu\, m^{\bar{n}_2}\big(f(\theta, \Delta) > 0\big), \quad \forall \theta \in \Theta, \forall \Gamma \in \mathcal{B}^{M\times m_1}, \tag{24}$$

for some $\mu \in (0, \infty)$, $\bar{n}_2 \geq n_2$ and $\epsilon \geq 0$. The SSP Algorithm is ϵ-secure for P_1 with confidence $1-\delta$, $\delta \in (0,1)$ as long as

$$n_1 = \left\lceil \frac{\mu}{\delta}(m_1 + 1) - 1 \right\rceil \bar{n}_2. \tag{25}$$

Additionally, suppose that we increase n_1 to satisfy

$$n_1 = \left\lceil \frac{\mu}{\delta}\Big(m_1 + \sqrt{2m_1 \ln\frac{1}{\beta}} + \ln\frac{1}{\beta}\Big)\right\rceil \bar{n}_2 \tag{26}$$

for some $\beta \in (0,1)$. Then, with probability larger than $1-\beta$, the SSP Algorithm generates a sampled security policy y^* with value $\bar{V}(A_1)$ that is ϵ-secure for P_1 with confidence $1-\delta$. □

Theorem 4.4 shows that, even when there is a mismatch in the distributions, it is still possible to achieve high-confidence security policies. However, the number of samples required by the SSP algorithm essentially needs to be multiplied by μ. Alternatively, if one uses the number of samples dictated by Theorem 3.1 and there is mismatch in the distributions, then one obtains security with confidence $1 - \delta/\mu$ (instead of $1-\delta$) since one can go from the formulas in Theorem 3.1 to the ones in Theorem 4.4 by simply replacing $1/\delta$ by μ/δ.

Remark 4.5 ($\mu < 1$) For values of $\epsilon > 0$ and matched (or closely matched distributions), the mismatch factors may actually be smaller than 1. Theorem 4.4 is still applicable and essentially states that if one is willing to accept some $\epsilon > 0$, one may get $1-\delta$ confidence with a smaller number of samples than what was required by Theorem 3.1. □

Remark 4.6 (Matrix-independent results) If we choose μ to satisfy

$$\tilde{m}^{\bar{n}_2}(B) \leq \mu\, m^{\bar{n}_2}(B), \quad \forall B \in \mathcal{F},$$

where $\mathcal{F}$ denotes the Borel sigma-algebra on $\mathcal{B}^{N\times\bar{n}_2}$, then (24) holds with $\epsilon = 0$ for every matrix game and we obtain a game independent result. The price is, of course, that such μ does not explore the structure of the particular game and may therefore be much larger than what is needed. In fact, we shall see in the next section that the structure of the matrix A may dictate that some mismatch should not lead to a degradation in the confidence levels. This is the case when A exhibits some form of policy domination. □

4.2.3. Matrix games with dominated policies

Consider a situation when P_1 knows of some particularly good policies that P_2 may apply to play the game. For example, suppose that the entries in some column $c_{\text{better-for-}\mathsf{P}_2}$ of A are all element-wise larger than those in some other column $c_{\text{worse-for-}\mathsf{P}_2}$. In this case, it turns out that P_1 can increase the probability of sampling the column $c_{\text{better-for-}\mathsf{P}_2}$ at the expense of decreasing the probability of selecting $c_{\text{worse-for-}\mathsf{P}_2}$ and this mismatch does not

[4] The confidence level $1-\beta$ refers to the extraction of $\Delta_1, \ldots \Delta_K$ that defines θ^*.

require a larger bound on the number of columns to sample. This observation is formalized in the remaining of this section.

We begin with the following notion of dominance.

Definition 1 (ϵ-Dominance) *Given an $M \times N$ matrix A, the vector $d^* \in \mathcal{B}^{N\times 1}$ is said to be* ϵ-dominated *by the vector $d \in \mathcal{B}^{N\times 1}$ for some $\epsilon \geq 0$ if*

$$e_i(M)'Ad^* \leq e_i(M)'Ad + \epsilon, \quad \forall i \in \{1, \ldots, M\},$$

where $e_i(M)$ denotes the ith element of the canonical basis of $\mathbb{R}^M$.

With $\epsilon = 0$, the above definition becomes identical to that of domination between pure policies in matrix games [cf. Basar and Olsder (1999)]. Next, we introduce the notion of two sampling distributions being perturbed.

Definition 2 (Perturbed sampling) *Given two distinct vectors $d, d^* \in \mathcal{B}^{N\times 1}$ and two probability measures $m, \tilde{m}$ on $\mathcal{B}^{N\times 1}$, we say that m is a* perturbation *of $\tilde{m}$ with respect to the pair (d, d^*) if*

1. *m differs from $\tilde{m}$ only over $\{d, d^*\} \subseteq \mathcal{B}^{N\times 1}$, i.e.,*

$$\tilde{m}(e_j(N)) = m(e_j(N)),$$

for all j such that $e_j(N) \notin \{d^*, d\}$, where $e_j(N)$ denotes the jth element of the canonical basis of $\mathbb{R}^N$;

2. *the probability of extracting d^* according to m is smaller than according to $\tilde{m}$, i.e.,*

$$m(d^*) \leq \tilde{m}(d^*).$$

We now present the main result of this subsection.

Theorem 4.7 (Domination) *Given the game matrix A, suppose that for some $\epsilon \geq 0$, there exist vectors $d^*, d \in \mathcal{B}^{N\times 1}$ such that d^* is ϵ-dominated by d. Suppose that Assumption 2.1 holds and that the columns of the matrices Π_1 and Π_2 are sampled according to distributions m and $\tilde{m}$, respectively. If m is a perturbation of $\tilde{m}$ with respect to (d, d^*), then Theorem 4.4 holds with $\mu = 1$.*

This result shows that even when P_1 extracts with low probability (possibly equal to zero) the column d^*, the bounds of Section 3 hold.

5. A-posteriori assessment of a given policy

Suppose now that P_1 obtained a policy y^* using either a randomized or a deterministic algorithm. In this section, we are interested in computing a high-confidence security level value $\bar{V}(y^*)$ for this policy, when y^* is played against P_2's policy z^* in (1). The *sampled security-value* (SSV) Algorithm 2 addresses this question.

Algorithm 2 [SSV Algorithm]

1: P_1 randomly selects k_1 columns of A, which corresponds to the selection of a random matrix $\bar{\Pi}_1 \in \mathcal{B}^{N\times k_1}$.

2: P_1 computes

$$\bar{V}(y^*) = \max_{j\in\{1,\ldots,k_1\}} y^{*\prime} A\bar{\Pi}_1 e_j(k_1), \tag{27}$$

where $e_j(k_1)$ denotes the jth element of the canonical basis of $\mathbb{R}^{k_1}$. We call $\bar{V}(y^*)$ P_1*'s a-posteriori sampled security value.*

We use the qualifier "a-posteriori" for the sampled security value $\bar{V}(y^*)$ to emphasize that this value is computed *after* a particular security policy y^* has been obtained. We say that *the SSV algorithm is ϵ-secure for player P_1's policy y^* with confidence* $1 - \delta$ if

$$\mathrm{P}_{\bar{\Pi}_1,\Pi_2}\left(y^{*\prime} Az^* \leq \bar{V}(y^*) + \epsilon \mid y^*\right) \geq 1 - \delta. \tag{28}$$

This condition states that the probability that the outcome of the game will violate P_1's a-posteriori sampled security value $\bar{V}(y^*)$ by more than ϵ is smaller than δ. As stated, this definition requires the bound to hold regardless of the algorithms used to generate y^* and z^*. In particular, we leave open the possibility that both policies could have been computed using the SSP algorithm, perhaps with confidence levels different than δ. In fact, one could imagine P_1 computing y^* using the SSP algorithm for a confidence level δ_{SSP} and then studying the security of such policy for tighter confidence levels δ using the SSV algorithm. Thus, the SSV algorithm combined with the SSP algorithm can be viewed as a heuristics for designing high-confidence security policies.

Also for the SSV, we can define a stronger notion of security that guarantees the inherent security of the a-posteriori sampled security value $\bar{V}(y^*)$, when P_1 plays y^* repeatedly against a sequence of policies z^* for P_2, each obtained by a distinct random exploration of her policy space. We say that *the a-posteriori sampled security value $\bar{V}(y^*)$ is ϵ-secure for player P_1's policy y^* with confidence* $1 - \delta$ if

$$\mathrm{P}_{\Pi_2}\left(y^{*\prime} Az^* \leq \bar{V}(y^*) + \epsilon \mid y^*,\ \bar{V}(y^*)\right). \tag{29}$$

As for the SSP Algorithm 1, we assume that the distributions of the column extraction matrices $\bar{\Pi}_1$ used by player P_1 in the SSV Algorithm 2 and Π_2 used by player P_2 in her randomized policy exploration are independent.

Assumption 5.1 (Independence) *The random matrix $\bar{\Pi}_1$ in the SSV Algorithm 2 and matrix Π_2 involved in player P_2 randomized exploration are statistically independent and each of them has independent and identically distributed columns.* □

We now provide a bound on the number of samples k_1 used in the SSV Algorithm to guarantee ϵ-security. Akin

to Section 3, we restrict our attention to the case when player P_1 uses the same distribution as player P_2 to extract columns of A. Results along the same lines as those shown in Section 4 for the SSP Algorithm could be obtain for the SSV Algorithm in the case of mismatched distributions, but we omit them because they are fundamentally similar.

Theorem 5.1 (SSV Algorithm) *Suppose that Assumption 5.1 holds and that $\bar{\Pi}_1 \in \mathcal{B}^{N\times k_1}$ and $\Pi_2 \in \mathcal{B}^{N\times n_2}$ have identically distributed columns. The SSV algorithm is $(\epsilon = 0)$-secure for player P_1's policy y^* with confidence $1-\delta$, $\delta \in (0,1)$ as long as*

$$k_1 = \left\lceil \frac{1}{\delta} - 1 \right\rceil \bar{n}_2, \tag{30}$$

with $\bar{n}_2 \geq n_2$. Additionally, suppose that we increase k_1 to satisfy

$$k_1 = \left\lceil \frac{1}{\delta} \ln \frac{1}{\beta} \right\rceil \bar{n}_2, \tag{31}$$

for some $\beta \in (0,1)$. Then, with probability $1-\beta$, the SSV Algorithm generates an a-posteriori sampled security value $\bar{V}(y^*)$ that is $(\epsilon = 0)$-secure for player P_1's policy y^* with confidence $1-\delta$. □

Proof of Theorem 5.1: Defining $K := \left\lceil \frac{1}{\delta} - 1 \right\rceil$ and the function $\bar{f} : \mathcal{S}_M \times \mathcal{B}^{N\times \bar{n}_2} \to \mathbb{R}$ by

$$\bar{f}(y, \Delta) = \max_{j\in\{1,\ldots,\bar{n}_2\}} y' A \Delta e_j(\bar{n}_2), \tag{32}$$

we can re-write (27) as

$$\bar{V}(y^*) = \max_{i\in\{1,\ldots,K\}} \bar{f}(y^*, \Delta_i), \tag{33}$$

where the matrices $\Delta_1, \Delta_2, \ldots, \Delta_K \in \mathcal{B}^{N\times\bar{n}_2}$ are obtained by partitioning the $K\bar{n}_2$ columns of $\bar{\Pi}_1 \in \mathcal{B}^{N\times K\bar{n}_2}$ into K i.i.d. matrices.

For any given y^* (which is independent of the Δ_i), we conclude from (Campi and Calafiore, 2009, Proposition 4) that the (conditional) probability that another matrix Δ, sampled independently from the same distribution as the Δ_i, satisfies the constraint

$$\bar{f}(y^*, \Delta) \leq \bar{V}(y^*) := \max_{i\in\{1,\ldots,K\}} \bar{f}(y^*, \Delta_i)$$

can be lower-bounded as follows:

$$\mathrm{P}_{\bar{\Pi}_1,\Delta}\left(\bar{f}(y^*,\Delta) \leq \bar{V}(y^*) \mid y^*\right) \geq \frac{K}{K+1} \geq 1-\delta, \tag{34}$$

where the second inequality is a consequence of (30). From the definition of $\bar{f}$, we conclude from (34) that

$$\mathrm{P}_{\bar{\Pi}_1,\Delta}(y^{*\prime} A\Delta e_j(\bar{n}_2) \leq \bar{V}(y^*), \forall j \in \{1,\ldots,\bar{n}_2\} \mid y^*) \geq 1-\delta,$$

and, since $n_2 \leq \bar{n}_2$, we also have that

$$\mathrm{P}_{\bar{\Pi}_1,\Delta}(y^{*\prime} A\Delta e_j(n_2) \leq \bar{V}(y^*), \forall j \in \{1,\ldots,n_2\} \mid y^*) \geq 1-\delta.$$

Under Assumption 5.1, when the columns of $\bar{\Pi}_1$ and Π_2 are identically distributed, the matrix consisting of the first n_2 columns of Δ can be viewed as the matrix Π_2 and the inequality above implies that

$$\mathrm{P}_{\bar{\Pi}_1,\Pi_2}(y^{*\prime} A\Pi_2 e_j(n_2) \leq \bar{V}(y^*), \forall j \in \{1,\ldots,n_2\} \mid y^*) \geq 1-\delta.$$

Since

$$y^{*\prime} A\Pi_2 e_j(n_2) \leq \bar{V}(y^*),\ \forall j \in \{1,\ldots,n_2\} \quad \Rightarrow \quad \Rightarrow \quad y^{*\prime} A\Pi_2 z \leq \bar{V}(y^*),\ \forall z \in \mathcal{S}_{n_2},$$

we conclude that

$$\mathrm{P}_{\Pi_2,\bar{\Pi}_1}\left(y^{*\prime} A\Pi_2 z_2^* \leq \bar{V}(y^*) \mid y^*\right) \geq 1-\delta,$$

which shows that SSV Algorithm is $(\epsilon = 0)$-secure with confidence $1-\delta$.

If, instead of using (Campi and Calafiore, 2009, Proposition 4) and (30) to obtain (34), we use (Campi and Garatti, 2008, Theorem 1) and (31), we obtain

$$\mathrm{P}_{\Delta}\left(\bar{f}(y^*,\Delta) \leq \bar{V}(y^*) \mid y^*, \bar{V}(y^*)\right) \geq 1-\delta, \tag{35}$$

with probability larger than $1-\beta$, where the confidence level $1-\beta$ refers to the extraction of $\bar{\Pi}_1 = [\Delta_1, \ldots, \Delta_K]$ that defines $\bar{V}(y^*)$. The proof can now proceed exactly as before, but with (34) replaced by (35), which now involves a probability conditioned to y^*, and $\bar{V}(y^*)$. This shows that if k_1 satisfies (31), then with probability larger than $1-\beta$, the security value $\bar{V}(y^*)$ is $(\epsilon = 0)$-secure with confidence $1-\delta$. ■

6. Example: Hide-and-seek matrix game

In this section, we apply the SSP and SSV Algorithms to a classic search problem: Consider a zero-sum game where P_1 hides a non-moving object (treasure) in one of N points $\{p_1, \ldots, p_N\} \subset \mathbb{R}^2$ on the plane and P_2 wants to find the treasure with minimum cost, by traveling from point to point until she finds it.

The game is played over the set of mixed policies:

- P_1 chooses a probability distribution $z \in \mathcal{S}_N$ for the treasure over the N points, and
- P_2 chooses a probability distribution $y \in \mathcal{S}_M$ over the set $\mathcal{R} := \{r_j : j = 1, \ldots, M\}$ of $M := N!$ routes that start at P_1's initial position $p_0 \in \mathbb{R}^2$ and go through all possible permutations of the points.

When P_1 chooses to hide the treasure at point p_i and P_2 selects route r_j, the outcome of the game is equal to the length of route r_j from P_1's initial position p_0 to the point p_i where the treasure lies. Namely,

$$A_{ij} = -\sum_{k=1}^{k_{ij}^*} \|r_j(k) - r_j(k-1)\|, \tag{36}$$

where $r_j(k) \in \mathbb{R}^2$, $k \in \{1, \ldots, N\}$ denotes the kth point in the route r_j with $r_j(0) = p_0$, and the summation ends at the index k^*_{ij} for which $r_j(k^*_{ij}) = p_i$ is the point where the treasure is hidden. The minus sign in (36) is needed to maintain consistency with the formulation in the first part of the paper, where P_1 is the minimizer. Indeed, P_1 hides the treasure to maximize the distance and therefore to minimize the entries of A.

For a large N, the exact computation of the optimal mixed strategies is intractable because the size of the matrix A is $N \times N!$. However, the results in this paper lead to a computational complexity that is *independent of the size of the game,* which means that we can provide probabilistic guarantees for games with an arbitrarily large number of points.

In this particular game, only the player P_2 that chooses paths has a large number of options ($M = N!$) so we can assume that both players consider all possible N locations where P_1 can hide the treasure (all rows of A), but randomly select only a small number of paths (columns of A) to construct their submatrices. However, the player P_1 that hides the treasure should respect the bounds provided by Theorems 3.1 and 5.1 to avoid unpleasant surprises.

In our numerical experiments, we considered $N = 10$ points distributed uniformly randomly in a square region. To illustrate the use of the SSP and the SSV Algorithms, we fixed $m_1 = N$, $\beta = 10^{-5}$, and $\bar{n}_2 = 10$ (Figure 1) or $\bar{n}_2 = 1000$ (Figure 2). To achieve a confidence level of $\delta = .01$ two approaches are possible:

SSP only: Execute the SSP Algorithm 1 with n_1 satisfying (6) to obtain a sampled security value and a sampled security policy with confidence $1-\delta = 99\%$.

SSP+SSV: Execute the SSP Algorithm 1 with a value for n_1 smaller than the one indicated by (6) to obtain a sampled security policy, and then run the SSV Algorithm 2 with a value of k_1 satisfying (31) to obtain an a-posteriori sampled security value with confidence $1 - \delta = 99\%$.

While the SSP+SSV option requires solving a smaller subgame in the SSP algorithm, and is therefore computationally more attractive, it typically results in a worst sampled security policy and therefore the corresponding security value is typically worst. However, one can see that the curves corresponding to the SSP+SSV option are relatively flat, which indicates that significant computational savings are possible without a significant degradation in the sampled security level. Note that the security levels computed using either of the approaches above are random variables since they depend on the randomly selected columns of the matrix A. The plots in Figures 1 and 2 show Monte Carlo estimates of the mean and standard deviation of these random variables.

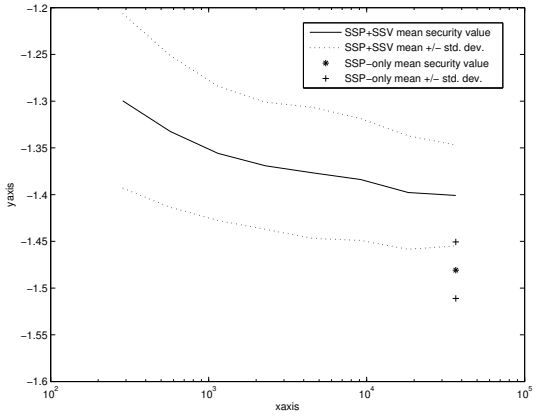

Figure 1: Numerically determined values for the 99% confidence sampled security value ($\delta = .1$). The solid line (mean) and dashed lines (plus/minus one standard deviation) were obtained using the SSP+SSV approach, using different values of n_1 in the SSP Algorithm (with n_1 in the x-axis) and a value for k_1 in the SSV Algorithm satisfying (31). The star '*' was obtained using the SSP-only approach, using the value for n_1 satisfying (6). The remaining parameters used are as follows: the number of points is $N = 10$, the side length of the square region is 1 unit, $m_1 = \bar{n}_2 = 10$, $\beta = 10^{-5}$, and the columns were drawn uniformly randomly. Each mean and standard deviation was estimated using 300 Monte Carlo samples.

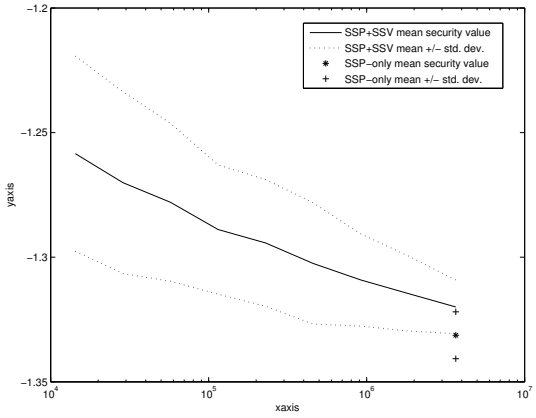

Figure 2: Plot similar to that in Figure 1, with the exception that we took $\bar{n}_2 = 1000$. Each mean and standard deviation was estimated using 300 Monte Carlo samples.

7. Conclusions and Future Directions

We addressed the solution of large zero-sum matrix games using randomized techniques. We provided a procedure based on randomized sampling by which a player can construct policies that are security with high-probability against an adversary engaged in a randomized exploration of the games characterized by large decision trees. We proposed a new probabilistic notion of security policy and level and derive bounds on the sample sizes that guarantees the discovery of a security policy with high probability. The bounds provided consider both the case where the two players sample policies using the same and different distributions. The applicability of the results is illustrated with a combinatorial hide-and-seek game.

This work suggests a number of future directions of research. One promising direction is to explore incremental opmization techniques to reduce the bound on the size of the submatrices and/or the number of entries of the submatrices that are needed to compute the sampled security policies. Another direction for future research regards the choice of the distributions used to sample policies to minimize the sample-size bounds and maximize the probability of finding adequate policies. In the context of the example in Section 6, we are currently exploring closed-loop versions of the hide-and-seek game that involve the searcher taking measurements regarding the location of the treasure as she moves from point to point [cf. Borri et al. (2011)].

References

Alamo, T., Tempo, R., Camacho, E. F., Nov. 2009. Randomized strategies for probabilistic solutions of uncertain feasibility and optimization problems. IEEE Trans. on Automat. Contr. 54 (11), 2545–2559.

Alamo, T., Tempo, R., Luque, A., Jun. 2010. On the sample complexity of randomized approaches to the analysis and design under uncertainty. In: Proc. of the 2010 Amer. Contr. Conf. Baltimore, MD, USA, pp. 4671–4676.

Basar, T., Olsder, G. J., 1999. Dynamic Non-Cooperative Game Theory. SIAM, Philadelphia, PA, USA.

Bellman, R., 1962. Dynamic programming treatment of the traveling salesman problem. J. Assoc. Comput. Mach. 9, 61–63.

Bopardikar, S. D., Borri, A., Hespanha, J. P., Prandini, M., Di Benedetto, M. D., Dec. 2010. Randomized sampling for large zero-sum games. In: Proc. of the 49th Conf. on Decision and Contr. Atlanta, GA, USA, pp. 7675–7680.

Bopardikar, S. D., Hespanha, J. P., Jun. 2011. Randomized solutions to partial information dynamic zero-sum games. In: Proc. of the 2011 Amer. Contr. Conf.

Borri, A., Bopardikar, S. D., Hespanha, J. P., Di Benedetto, M. D., Aug. 2011. Hide-and-seek with directional sensing. In: Proc. of the 18th World Congress of Int. Federation of Automat. Contr. Milan, Italy.

Browne, C., Powley, E., Whitehouse, D., Lucas, S., Cowling, P., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., Colton, S., March 2012. A survey of Monte Carlo tree search methods. IEEE Transactions on Computational Intelligence and AI in Games 4 (1), 1–43.

Calafiore, G., 2009. On the expected probability of constraint violation in sampled convex programs. Journal of Optimization Theory and Applications 143 (2), 405–412.

Calafiore, G., 2010. Random convex programs. SIAM Journal of Optimization 20 (6), 3427–3463.

Calafiore, G. C., Campi, M. C., May 2006. The scenario approach to robust control design. IEEE Trans. on Automat. Contr. 51 (5), 742–753.

Campi, M. C., Calafiore, G. C., Feb. 2009. Notes on the scenario design approach. IEEE Trans. on Automat. Contr. 54 (2), 382–385.

Campi, M. C., Garatti, S., 2008. The exact feasibility of randomized solutions of robust convex programs. SIAM J. Contr. Optimization 19 (3), 1211–1230.

Campi, M. C., Garatti, S., Prandini, M., Dec. 2009. The scenario approach for systems and control design. Annual Reviews in Control 33 (2), 149–157.

de Farias, D. P., Roy, B. V., Aug. 2004. On constraint sampling in the linear programming approach to approximate dynamic programming. Mathematics of Operations Research 29 (3), 462–478.

Erdoğan, E., Iyengar, G., Dec. 2006. Ambiguous chance constrained problems and robust optimization. Mathematical Programming 107 (1–2), 37–61.

Frank, A., 1989. Heuristic programming in Artificial Intelligence 2: The second computer olympiad. Ellis Horwood, Ch. Brute force search in games of imperfect information, pp. 204–209.

Frank, I., Basin, D., 1999. Optimal play against best defence: Complexity and heuristics. In: Herik, H., Iida, H. (Eds.), Computers and Games. Vol. 1558 of Lecture Notes in Computer Science. Springer Berlin Heidelberg, pp. 50–73.

Fujisaki, Y., Kozawa, Y., December 2003. Probabilistic robust controller design: Probable near-minimax value and randomized algorithm. In: Proceedings of the IEEE Conference on Decision and Control. Maui, Hawaii, USA, pp. 1938–1943.

Ginsberg, M., 1996. Partition search. In: Thirteenth National Conference on Artificial Intelligence (AAAI-96). pp. 228–233.

Hespanha, J. P., Prandini, M., Dec. 2001. Nash equilibria in partial-information games on Markov chains. In: Proc. of the 40th Conf. on Decision and Contr. pp. 2102–2107.

Hinton, A., Kwiatkowska, M., Norman, G., Parker, D., 2006. PRISM: A tool for automatic verification of probabilistic systems. In: Hermanns, H., Palsberg, J. (Eds.), Tools and Algorithms for the Construction and Analysis of Systems. Vol. 3920 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, pp. 441–444.

Khargonekar, P., Tikku, A., Dec. 1996. Randomized algorithms for robust control analysis and synthesis have polynomial complexity. In: Proc. of the 35th Conf. on Decision and Contr. Vol. 3. Kobe, Japan, pp. 3470–3475.

LaValle, S. M., Kuffner, J. J., 2001. Rapidly-exploring random trees: Progress and prospects. In: Donald, B. R., Lynch, K. M., , Rus, D. (Eds.), Algorithmic and Computational Robotics: New Directions. Wellesley, MA, pp. 293–308.

Lipton, R. J., Young, N. E., 1994. Simple strategies for large zero-sum games with applications to complexity theory. In: Twenty-sixth annual ACM Symposium on Theory of Computing. pp. 734–740.

Lye, K.-W., Wing, J., 2005. Game strategies in network security. International Journal of Information Security 4, 71–86.

Motwani, R., Raghavan, P., 1995. Randomized Algorithms. Cambridge University Press.

Parker, A., Nau, D., Subrahmanian, V. S., 2005. Game-tree search with combinatorially large belief states. In: International Joint Conference on Artificial Intelligence. pp. 254–259.

Tempo, R., Bai, E. W., Dabbene, F., 1997. Probabilistic robustness analysis: Explicit bounds for the minimum number of samples. Syst. & Contr. Lett. 30 (5), 237–242.

Tempo, R., Calafiore, G., Dabbene, F., 2004. Randomized Algorithms for Analysis and Control of Uncertain Systems. Springer-Verlag, London.

Vapnik, V., 1998. Statistical Learning Theory. John Wiley, New York.

Vidyasagar, M., Dec. 1998. Statistical learning theory and randomized algorithms for control. IEEE Contr. Syst. Mag. 18 (6), 69–85.

Vidyasagar, M., Blondel, V. D., Sep. 2001. Probabilistic solutions to some NP–hard matrix problems. Automatica 37 (9), 1397–1405.

Von Neumann, J., 1928. Zur theorie der gesellschaftsspiele. Math. Annalen. 100, 295–320.

Appendix

In this Appendix, we provide the proofs of the results in Section 4.

Proof of Theorem 4.1: Following the same steps as in the proof of Theorem 3.1, we can conclude that the sampled security value $\bar{V}(A_1)$ can be expressed as:

$$\bar{V}(A_1) = \min_{\theta \in \Theta} \Big\{ v : f_{\Gamma_1}(\theta, \Delta_i) \leq 0 \; \forall i \in \{1, \ldots, K\} \Big\}, \quad (37)$$

where $\theta = (y, v) \in \mathcal{S}_{m_1} \times \mathbb{R}$, f_{Γ_1} is defined in (8), and

$$K := \Big\lceil 2(m_1+1) + \frac{2}{\delta-\rho}\ln\frac{1}{\beta} + \frac{2(m_1+1)}{\delta-\rho}\ln\frac{2}{\delta-\rho}\Big\rceil, \tag{38}$$

and $\Delta_1, \Delta_2, \ldots, \Delta_K$ are i.i.d. matrices obtained from the columns of $\Pi_1 \in \mathcal{B}^{N\times K\bar{n}_2}$ with probability distribution $m^{\bar{n}_2}$. Since $\rho \in [0,1)$,

$$\Delta = \Delta_i \;\Leftrightarrow\; d(\Delta,\Delta_i) \le \rho, \quad \forall \Delta, \Delta_i \in \mathcal{B}^{N\times\bar{n}_2}$$

where d denotes the discrete metric (16). Therefore (37) is equivalent to

$$\bar{V}(A_1) = \min_{\theta\in\Theta} \big\{ v : f_{\Gamma_1}(\theta,\Delta) \le 0,\, \forall\, \Delta \text{ such that } d(\Delta,\Delta_i) \le \rho \text{ for some } i \in \{1,\ldots,K\}\big\}. \tag{39}$$

For an arbitrary realization of the matrix Γ_1 (which is independent of the Δ_i), we can conclude from (Erdoğan and Iyengar, 2006, Theorem 6) that, for any random variable Δ with measure $\tilde{m}^{\bar{n}_2}$, we have that

$$\tilde{m}^{\bar{n}_2}\big(f_{\Gamma_1}(\theta^*,\Delta) \le 0 \mid \Gamma_1, \theta^*\big) \ge 1-\delta, \tag{40}$$

with probability at least

$$1 - \Big(\frac{\mathrm{e}K}{m_1+1}\Big)^{m_1+1} \mathrm{e}^{-(\delta-\rho)(K-(m_1+1))}, \tag{41}$$

where θ^* denotes the value in Θ that achieves the minimum in (39). Here, the confidence refers to the extraction of matrix $\Pi_1 = [\Delta_1, \ldots, \Delta_K]$ that defines θ^* (given Γ_1). Since the minimum in (39) is achieved for $\theta^* = (y_1^*, \bar{V}(A_1))$, we can use the definition (8) of f_{Γ_1} to re-write (40) as

$$\tilde{m}^{\bar{n}_2}\big(y_1^{*\prime}\Gamma_1' A\Delta e_j(\bar{n}_2) \le \bar{V}(A_1),\ \forall j \in \{1,\ldots,\bar{n}_2\} \mid \Gamma_1, y_1^*, \bar{V}(A_1)\big) \ge 1-\delta.$$

Under Assumption 2.1, the matrix consisting of the first n_2 columns of Δ can be viewed as the matrix Π_2 and the inequality above implies that

$$\mathrm{P}_{\Pi_2}\big(y_1^{*\prime}\Gamma_1' A\Pi_2 e_j(n_2) \le \bar{V}(A_1),\ \forall j \in \{1,\ldots,n_2\} \mid \Gamma_1, y_1^*, \bar{V}(A_1)\big) \ge 1-\delta.$$

Since

$$y_1^{*\prime}\Gamma_1' A\Pi_2 e_j(n_2) \le \bar{V}(A_1),\ \forall j \in \{1,\ldots,n_2\} \;\Rightarrow\; y_1^{*\prime}\Gamma_1' A\Pi_2 z \le \bar{V}(A_1),\ \forall z \in \mathcal{S}_{n_2},$$

we conclude that

$$\mathrm{P}_{\Gamma_2,\Pi_2}\big(y_1^{*\prime}\Gamma_1' A\Pi_2 z_2^* \le \bar{V}(A_1) \mid \Gamma_1, y_1^*, \bar{V}(A_1)\big) \ge 1-\delta.$$

Since we have shown that this bound holds for an arbitrary realization of Γ_1, it also holds for the unconditional probability. This shows that, with probability at least (41), the policy y^* with value $\bar{V}(A_1)$ is ($\epsilon = 0$)-secure for P_1 with confidence $1-\delta$. To conclude the proof, we only need to show that (38) implies that the probability (41) is larger than $1-\beta$. Defining, $m := m_1 + 1$ and $\bar{\delta} := \delta - \rho$, we conclude from (38) that

$$\frac{K}{2} \ge m - \frac{1}{\bar{\delta}}\ln\beta + \frac{m}{\bar{\delta}}\ln\frac{2}{\bar{\delta}},$$

which implies that

$$\begin{aligned} K &\ge m - \frac{1}{\bar{\delta}}\ln\beta + \frac{m}{\bar{\delta}}\ln\frac{2}{\bar{\delta}} + \frac{K}{2} \\ &= m - \frac{1}{\bar{\delta}}\ln\beta + \frac{m}{\bar{\delta}}\Big(\ln\frac{2}{\bar{\delta}} + \frac{\bar{\delta}K}{2m}\Big) \\ &\ge m - \frac{1}{\bar{\delta}}\ln\beta + \frac{m}{\bar{\delta}}\Big(\ln\frac{2}{\bar{\delta}} + 1 - \ln\frac{2m}{\bar{\delta}K}\Big) \\ &= m - \frac{1}{\bar{\delta}}\ln\beta + \frac{m}{\bar{\delta}}\ln\frac{eK}{m}, \end{aligned}$$

where the second inequality is a consequence of the fact that $\frac{1}{x} \ge 1 - \ln x$, $\forall x > 0$. Therefore

$$\ln\beta \ge -\bar{\delta}(K-m) + m\ln\frac{eK}{m} \Leftrightarrow \beta \ge \Big(\frac{\mathrm{e}K}{m}\Big)^m \mathrm{e}^{-\bar{\delta}(K-m)},$$

which confirms that (41) is larger than $1-\beta$. ■

Proof of Lemma 4.2: Given that

$$\mathrm{P}_{\bar{\Delta},\Delta_1,\ldots,\Delta_K}\big(f(\theta^*,\bar{\Delta}) > \epsilon\big) = \mathbb{E}_{\Delta_1,\ldots,\Delta_K}\Big[\mathrm{P}_{\bar{\Delta}}\big(f(\theta^*,\bar{\Delta}) > \epsilon \big| \theta^*\big)\Big],$$

we can use (20) to conclude that

$$\begin{aligned} \mathrm{P}_{\bar{\Delta},\Delta_1,\ldots,\Delta_K}\big(f(\theta^*,\bar{\Delta}) > \epsilon\big) &\le \mu\mathbb{E}_{\Delta_1,\ldots,\Delta_K}\Big[\mathrm{P}_{\Delta}\big(f(\theta^*,\Delta) > 0 \big| \theta^*\big)\Big] \\ &\le \frac{\mu\, n_\theta}{K+1}, \end{aligned}$$

where the second inequality follows from (Campi and Calafiore, 2009, Proposition 3). ■

Proof of Lemma 4.3: By (Alamo et al., 2010, Theorem 4), if we fix an arbitrary $\bar{\delta} \in (0,1)$ and K satisfies

$$K \ge \Big\lceil \frac{1}{\bar{\delta}}\Big(\ln\frac{1}{\beta} + (n_\theta - 1) + \sqrt{2(n_\theta-1)\ln\frac{1}{\beta}}\Big)\Big\rceil,$$

then, with probability [5] larger than $1-\beta$,

$$\mathrm{P}_{\Delta}\big(f(\theta^*,\Delta) > 0 \,|\, \theta^*\big) \le \bar{\delta}.$$

From definition of $\mu_f(\epsilon)$ in (20) it follows that

$$\mathrm{P}_{\bar{\Delta}}\big(f(\theta^*,\bar{\Delta}) > \epsilon \,|\, \theta^*\big) \le \mu_f(\epsilon)\,\mathrm{P}_{\Delta}\big(f(\theta^*,\Delta) > 0 \,|\, \theta^*\big).$$

[5] The confidence level $1-\beta$ refers to the extraction of $\Delta_1, \ldots \Delta_K$ that defines θ^*.

By combining the last two inequalities, we obtain

$$\begin{aligned}\mathrm{P}_{\bar{\Delta}}\left(f(\theta^*,\bar{\Delta})\le\epsilon\,|\,\theta^*\right) &= 1-\mathrm{P}_{\bar{\Delta}}\left(f(\theta^*,\bar{\Delta})>\epsilon\,|\,\theta^*\right)\\ &\ge 1-\mu_f(\epsilon)\bar{\delta},\end{aligned}$$

so that setting $\bar{\delta}:=\delta/\mu_f(\epsilon)$, the claim is proved. ■

Proof of Theorem 4.4: Since matrix Π_1 has $n_1 = K\bar{n}_2$ columns with $K=\left\lceil \mu\frac{m_1+1}{\delta}-1\right\rceil$, following the same steps as in the proof of Theorem 3.1, we can partition Π_1 into K i.i.d. matrices $\Delta_1,\Delta_2,\ldots,\Delta_K$, each in the set $\mathcal{B}^{N\times\bar{n}_2}$, and express the security value $\bar{V}(A_1)$ as:

$$\bar{V}(A_1)=\min_{\theta\in\Theta}\left\{v: f_{\Gamma_1}(\theta,\Delta_i)\le 0\ \forall i\in\{1,\ldots,K\}\right\}, \quad (42)$$

with $\theta=(y,v)\in\Theta=\mathcal{S}_{m_1}\times\mathbb{R}$, and f_Γ defined in (23).

Let the minimum in (42) be achieved for some $\theta^*=(y_1^*,\bar{V}(A_1))$. Matrices Δ_i are random variables distributed according to $\mathrm{P}_\Delta=m^{\bar{n}_2}$ over $\mathcal{B}^{N\times\bar{n}_2}$. For any given realization of the matrix Γ_1 (which is independent of the Δ_i by Assumption 2.1) we conclude from Lemma 4.2 and the definition of μ in (24), that the (conditional) probability that another matrix $\bar{\Delta}$, sampled independently of the Δ_i according to probability $\mathrm{P}_{\bar{\Delta}}=\tilde{m}^{\bar{n}_2}$ over $\mathcal{B}^{N\times\bar{n}_2}$, satisfies

$$\mathrm{P}_{\Pi_1,\bar{\Delta}}\left(f_{\Gamma_1}(\theta^*,\bar{\Delta})\le\epsilon\mid\Gamma_1\right)\ge\frac{K-m_1}{K+1}\ge 1-\delta, \quad (43)$$

where the second inequality is a consequence of (25). Using the definition of f_{Γ_1} from (8) and θ^*, we can re-write (43) as

$$\begin{aligned}\mathrm{P}_{\Pi_1,\bar{\Delta}}\big(&y_1^{*\prime}\Gamma_1'A\bar{\Delta}e_j(\bar{n}_2)\le\bar{V}(A_1)+\epsilon,\\ &\forall j\in\{1,\ldots,\bar{n}_2\}\mid\Gamma_1\big)\ge 1-\delta.\end{aligned}$$

Since $n_2\le\bar{n}_2$, we further conclude that

$$\begin{aligned}\mathrm{P}_{\Pi_1,\bar{\Delta}}\big(&y_1^{*\prime}\Gamma_1'A\bar{\Delta}e_j(n_2)\le\bar{V}(A_1)+\epsilon,\\ &\forall j\in\{1,\ldots,n_2\}\mid\Gamma_1\big)\ge 1-\delta.\end{aligned}$$

Under Assumption 2.1, the matrix consisting of the first n_2 columns of $\bar{\Delta}$ can be viewed as the matrix Π_2 and we conclude from the inequality above that

$$\begin{aligned}\mathrm{P}_{\Pi_1,\Pi_2}\big(&y_1^{*\prime}\Gamma_1'A\Pi_2e_j(n_2)\le\bar{V}(y^*),\\ &\forall j\in\{1,\ldots,n_2\}\mid\Gamma_1\big)\ge 1-\delta.\end{aligned}$$

Since

$$\begin{aligned}y_1^{*\prime}\Gamma_1'A\Pi_2e_j(n_2)\le\bar{V}(y^*),\ \forall j\in\{1,\ldots,n_2\}\quad\Rightarrow\\ \Rightarrow\quad y_1^{*\prime}\Gamma_1'A\Pi_2z\le\bar{V}(y^*),\ \forall z\in\mathcal{S}_{n_2},\end{aligned}$$

we conclude that

$$\mathrm{P}_{\Pi_1,\Gamma_2,\Pi_2}\left(y_1^{*\prime}\Gamma_1'A\Pi_2z_2^*\le\bar{V}(y^*)\mid\Gamma_1\right)\ge 1-\delta.$$

Since we have shown that this bound holds for an arbitrary realization of Γ_1, it also holds for the unconditional probability, which shows that the SSP algorithm is ϵ-secure for P_1 with confidence $1-\delta$.

If instead of applying Lemma 4.2, we apply Lemma 4.3, then using (26), we conclude that

$$\mathrm{P}_{\bar{\Delta}}\left(f_{\Gamma_1}(\theta^*,\bar{\Delta})\le\epsilon\mid\Gamma_1,\theta^*\right)\ge 1-\delta$$

with probability larger than $1-\beta$, where the confidence level $1-\beta$ refers to the extraction of $\Pi_1=[\Delta_1,\ldots,\Delta_K]$ that defines θ^* (given Γ_1). The proof can now proceed exactly as before, but with (43) replaced by the inequality above, which now involves a probability conditioned to y^* and $\bar{V}(A_1)$. This shows that, with probability larger than $1-\beta$, the policy y^* with value $\bar{V}(A_1)$ is ϵ-secure for P_1 with confidence $1-\delta$.

■

Proof of Theorem 4.7: To prove Theorem 4.7, we just need to show that

$$\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta})>\epsilon\right)\le\mathrm{P}_\Delta\left(f_\Gamma(\theta,\Delta)>0\right), \quad (44)$$

for any $\theta\in\Theta$, $\Gamma\in\mathcal{B}^{M\times m_1}$, since from this condition we have that $\mu=1$ satisfies (24).

Fix $\theta=(y_1,v)\in\Theta$ and $\Gamma\in\mathcal{B}^{M\times m_1}$. Let us distinguish between the following two cases:
Case 1) d^* satisfies

$$y_1'\Gamma'Ad^*-v>\epsilon \quad (45)$$

Case 2) d^* satisfies

$$y_1'\Gamma'Ad^*-v\le\epsilon \quad (46)$$

Starting with Case 1, observe that

$$\mathrm{P}_\Delta\left(f_\Gamma(\theta,\Delta)>0\right)=1-\mathrm{P}_\Delta\left(f_\Gamma(\theta,\Delta)\le 0\right). \quad (47)$$

Now, given that the columns of Δ are extracted independently according to m, we have

$$\mathrm{P}_\Delta\left(f_\Gamma(\theta,\Delta)\le 0\right)=\sum_{\{\Delta=[c_1,\ldots,c_{\bar{n}_2}]:f_\Gamma(\theta,\Delta)\le 0\}}\prod_{j=1}^{\bar{n}_2}m(c_j),$$

where $c_j\in\mathcal{B}^{N\times 1}$ denotes the jth column of Δ. Since from Definition 1 of ϵ-dominance and equation (45) we obtain

$$y_1'\Gamma'Ad-v>y_1'\Gamma'Ad^*-v-\epsilon>0,$$

we can conclude that the columns c_j, $j=1,\ldots,\bar{n}_2$, of Δ such that $f_\Gamma(\theta,\Delta)=\max_{j\in\{1,\ldots,\bar{n}_2\}}y_1'\Gamma'A\Delta e_j(\bar{n}_2)-v=\max_{j\in\{1,\ldots,\bar{n}_2\}}y_1'\Gamma'Ac_j-v\le 0$ must be different from both d and d^*. By Definition 2, we then have that $m(c_j)=\tilde{m}(c_j)$, $j=1,\ldots,\bar{n}_2$, and, hence,

$$\begin{aligned}\mathrm{P}_\Delta\left(f_\Gamma(\theta,\Delta)\le 0\right)&=\sum_{\{\Delta=[c_1,\ldots,c_{\bar{n}_2}]:f_\Gamma(\theta,\Delta)\le 0\}}\prod_{i=1}^{\bar{n}_2}\tilde{m}(c_i)\\ &=\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta})\le 0\right).\end{aligned}$$

Now, if we use this in equation (47), we get that

$$\begin{aligned}\mathrm{P}_\Delta\left(f_\Gamma(\theta,\Delta)>0\right)&=\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta})>0\right)\\ &\ge\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta})>\epsilon\right),\end{aligned}$$

i.e., equation (44) holds in Case 1.

Regarding Case 2, we start considering the case when d satisfies

$$y_1'\Gamma' Ad - v > \epsilon.$$

Note that $\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta}) \le \epsilon\right)$ can be expressed as follows

$$\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta}) \le \epsilon\right) = \sum_{\{\bar{\Delta}=[\bar{c}_1,\ldots,\,\bar{c}_{\bar{n}_2}]:\, f_\Gamma(\theta,\bar{\Delta})\le\epsilon\}} \prod_{j=1}^{\bar{n}_2} \bar{m}(\bar{c}_j)$$

Since the columns $\bar{c}_j$, $j = 1,\ldots,\bar{n}_2$, of any $\bar{\Delta} = [\bar{c}_1,\ldots,\,\bar{c}_{\bar{n}_2}]$ such that $f_\Gamma(\theta,\bar{\Delta}) = \max_{j\in\{1,\ldots,\bar{n}_2\}} y_1'\Gamma' A\bar{c}_j - v \le \epsilon$ must be different from d and the probability of extracting any such column according to $\bar{m}$ is larger than according to m (see Definition 2), we get

$$\begin{aligned}\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta}) \le \epsilon\right) &\ge \sum_{\{\bar{\Delta}=[\bar{c}_1,\ldots,\,\bar{c}_{\bar{n}_2}]:\, f_\Gamma(\theta,\bar{\Delta})\le\epsilon\}} \prod_{j=1}^{\bar{n}_2} m(\bar{c}_j)\\ &= \mathrm{P}_{\Delta}\left(f_\Gamma(\theta,\Delta) \le \epsilon\right)\end{aligned}$$

From this it follows that

$$\begin{aligned}\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta}) > \epsilon\right) &= 1 - \mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta}) \le \epsilon\right)\\ &\le 1 - \mathrm{P}_{\Delta}\left(f_\Gamma(\theta,\Delta) \le \epsilon\right)\\ &= \mathrm{P}_{\Delta}\left(f_\Gamma(\theta,\Delta) > \epsilon\right)\\ &\le \mathrm{P}_{\Delta}\left(f_\Gamma(\theta,\Delta) > 0\right),\end{aligned}$$

i.e. equation (44) holds.

We shall consider now the last subcase when d satisfies

$$y_1'\Gamma' Ad - v \le \epsilon. \tag{48}$$

We start noting that $\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta}) > \epsilon\right)$ is the probability that at least one of the columns, say $\bar{c}$, of $\bar{\Delta}$ satisfies

$$y_1'\Gamma' A\bar{c} - v > \epsilon.$$

Let

$$C = \{c \in \mathcal{B}^{N\times 1} : y_1'\Gamma' Ac - v > \epsilon\}.$$

Set $p_C = \sum_{c\in C} m(c)$ and $\tilde{p}_C = \sum_{c\in C} \tilde{m}(c)$. Then,

$$\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta}) > \epsilon\right) = 1 - (1-\tilde{p}_C)^{\bar{n}_2},$$

where $(1-\tilde{p}_C)^{\bar{n}_2}$ is the probability of all $\bar{n}_2$ independently extracted columns of $\bar{\Delta}$ not belonging to set C. Similarly,

$$\mathrm{P}_{\Delta}\left(f_\Gamma(\theta,\Delta) > \epsilon\right) = 1 - (1-p_C)^{\bar{n}_2}.$$

Now, since $C \cap \{d, d^*\} = \emptyset$ (see equations (46) and (48) and the definition of set C), from Definition 2 it follows that $p_C = \tilde{p}_C$, and therefore

$$\begin{aligned}\mathrm{P}_{\bar{\Delta}}\left(f_\Gamma(\theta,\bar{\Delta}) > \epsilon\right) &= \mathrm{P}_{\Delta}\left(f_\Gamma(\theta,\Delta) > \epsilon\right)\\ &\le \mathrm{P}_{\Delta}\left(f_\Gamma(\theta,\Delta) > 0\right).\end{aligned}$$

Given that we have shown that equation (44) holds for arbitrary values of θ and Γ, the proof is completed. ■