

Automated Generation of Counter-Terrorism Policies using Multi-Expert Input

ANSHUL SAWANT, University of Maryland
 JOHN P. DICKERSON, Carnegie Mellon University
 MOHAMMAD T. HAJIAGHAYI, University of Maryland
 V.S. SUBRAHMANYAN, University of Maryland

The use of game theory to model conflict has been studied by several researchers, spearheaded by Schelling [1960]. Most of these efforts assume a single payoff matrix that captures players' utilities under different assumptions about what the players will do. Our experience in counter-terrorism applications is that experts disagree on these payoffs. We leverage Shapley's [1959] notion of vector equilibria, which formulates games where there are multiple payoff matrices, but note that they are very hard to compute in practice. In order to effectively enumerate large numbers of equilibria with payoffs provided by multiple experts, we propose a novel combination of vector payoffs and well-supported ϵ -approximate equilibria [Daskalakis et al. 2006b]. We develop bounds related to computation of these equilibria for some special cases, and give a quasi-polynomial time approximation scheme (QPTAS) for the general case when the number of players is small (which is true in many real-world applications). Leveraging this QPTAS, we give efficient algorithms to find such equilibria and experimental results showing they work well on simulated data.

We then built a policy recommendation engine based on vector equilibria, called PREVE. We use PREVE to model the terrorist group Lashkar-e-Taiba (LeT), responsible for the 2008 Mumbai attacks, as a five-player game. Specifically, we apply it to three payoff matrices provided by experts in India-Pakistan relations, analyze the equilibria generated by PREVE, and suggest counter-terrorism policies that may reduce attacks by LeT. We briefly discuss these results and identify their strengths and weaknesses from a policy point of view.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics; I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms: Algorithms, Economics, Experimentation

Additional Key Words and Phrases: Counter-terrorism, game theory, vector equilibria, behavioral modeling, Lashkar-e-Taiba

ACM Reference Format:

Anshul Sawant, John P. Dickerson, Mohammad T. Hajiaghayi, and V.S. Subrahmanian. 2014. Automated Generation of Counter-Terrorism Policies using Multi-Expert Input. *ACM Trans. Intell. Syst. Technol.* 1, 1, Article 1 (September 2014), 27 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

This work is supported by the US Army Research Office under grants W911NF0910206 and W911NF1110344, NSF grant SES0826886, NSF CAREER award 1053605, NSF grant CCF-1161626, ONR YIP award N000141110662, and DARPA/AFOSR grant FA9550-12-1-0423. Dickerson was supported by a National Defense Science & Engineering (NDSEG) fellowship through the Army Research Office.

Authors' addresses: Anshul Sawant, University of Maryland, College Park, MD 20742; email: asawant@cs.umd.edu; John P. Dickerson, Carnegie Mellon University, Pittsburgh, PA 15213; email: dickerson@cs.cmu.edu; Mohammad T. Hajiaghayi, University of Maryland, College Park, MD 20742; email: hajiagha@cs.umd.edu; V.S. Subrahmanian, University of Maryland, College Park, MD 20742; email: vs@cs.umd.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2014 ACM 2157-6904/2014/09-ART1 \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

The research reported in this paper was motivated by a concrete application: how can countries trying to rein in the terrorist group Lashkar-e-Taiba¹ (LeT for short) come up with policies against them, especially if these policies need to be coordinated? In the case of a five-player game we formulated for LeT (presented later), there were wide variations of opinion amongst experts on what to do about LeT with respect to, for instance, whether India should carry out covert action, carry out coercive diplomacy, propose peace talks, or just keep the status quo. Likewise, the US has historically had multiple opposing viewpoints on whether to continue financial (development and military aid) to Pakistan, whether to carry out covert action against LeT, or do nothing. Analyzing the benefits of these actions even in the case of a single actor (e.g., only India or only the US) has proven challenging. The main contribution of this paper is a multi-player, game-theoretic framework in which this specific problem can be solved.

However, we wanted to come up with a general solution, one that is applicable to many different settings. For instance, there are many applications where the “payoff matrices”, usually one of the very first things needed in any game-theoretic framework, cannot be specified with accuracy. When asked about payoffs, multiple experts might express substantial disagreements. This is what happened with our LeT application. Here are some applications where multiple payoff matrices have been considered in a wide variety of settings.

- (1) *Socio-Cultural Behavior Modeling*. Woodley et al. [2008] propose a “Culturally Aware Response” (or CAR) framework in conjunction with the well-known World Values Survey to assess the results of different types of interactions between culturally different groups. They use multiple payoff matrices in their framework which vary based on the historical behaviors of different groups, e.g., one payoff matrix may indicate situations where a player responds in kind to responses of other players, while another payoff matrix may reflect situations where the player is largely non-violent.
- (2) *Open Source Software Releases*. Asundi et al. [2012] analyze the circumstances that are optimal for companies to release software. They argue that by open-sourcing a “crimped” version of their product, a company can hurt competitors, while enabling sales of a more sophisticated pay version of their product. To build their model, the authors utilize four different payoff matrices, corresponding to different regions of the parameter space that defines their model.
- (3) *International Climate Change Negotiations*. Pittel and Rübhelke [2012] develop a game-theoretic model of climate change negotiations building upon the well-known chicken game and the iterated prisoner’s dilemma. The two games are combined into a 3×3 game and studied under different payoff scenarios.
- (4) *Telecommunications*. Karami and Glisic [2010] define *asymmetric matrix games* (AMG) with which they model routing and network coding using conflict-free scheduling mechanisms. In their framework, multiple payoff matrices are defined, with one payoff matrix corresponding to each of a set of different partial possible network topologies.

Other applications include international negotiations [Kraus et al. 1995] where the precise payoffs for the nations involved are viewed through different lenses by different experts. They can also include applications where there are different views on

¹Lashkar-e-Taiba (LeT), translated variously from Urdu into “Army of the Pure” or “Army of the Pious”, is a prominent south Asian terrorist organization responsible for attacks in India, Kashmir, Pakistan, and Afghanistan, including the three days of attacks in 2008 in Mumbai, India, that resulted in the deaths of 166 innocent people [Tankel 2011a; Subrahmanian et al. 2012].

Table I. The actions that different players can take.

Player	Action	Abbrev.
Lashkar-e-Taiba (LeT)	Launch major attacks	attack
	Eliminate armed wing	eaw
	Hold attacks	hold
	Do nothing	none
Pakistan's Government (PakG)	Prosecute LeT	pros
	Endorse LeT	endorse
	Do nothing	none
Pakistan's Military (PakM)	Crackdown on LeT	crack
	Cut support to LeT	cut
	Increase support to LeT	support
	Do nothing	none
India	Covert action against LeT	covert
	Coercive diplomacy against PakG	coerce
	Propose peace initiative to PakG	peace
	Do nothing	none
U.S.	Covert action against LeT	covert
	Cut aid to PakG	cut
	Expand aid to PakG	expand
	Do nothing	none

the payoffs different corporations get for taking different types of actions (e.g., raise wages for striking workers vs. shut down a factory vs. take legal action). Even a seemingly simple action such as “take legal action” can lead to a diversity of views about costs/payoffs as different views may exist on, e.g., how long the litigation will take (and hence how much it will cost). Our paper has two parts:

- *Approximate Equilibria for Multi-Player Games with Vector Payoffs.* Games with multiple payoffs were introduced by Shapley [1959]. Shapley called them vector valued games and they have been extensively studied under various other names such as multicriteria games and multi-objective games. Unfortunately, for real-world applications such as the LeT application motivating this research, the computational cost of these past methods is too high. In order to address this, we introduce a novel combination of vector valued games and approximate equilibria and define new types of approximate equilibria for games with multiple players and multiple payoff matrices. We design algorithms for computing such equilibria for zero-sum games and games of low rank. For the case of rank 1 games, we give a structural result and use it to design a simple algorithm for such games. For general games we give an extension of Althöfer’s Approximation Lemma [1994] for simultaneous games with multiple payoff functions (SGMs) and use it to design a quasi-polynomial time approximation scheme (QPTAS) when the number of players in a game is constant (which is the case for our LeT game).
- *Application of PREVE to Generate Policies to Reduce Terror Acts by LeT.* Building on work by Dickerson et al. [2011; 2013], we then present a real-world application in which there are five parties including four governmental entities and the terrorist group Lashkar-e-Taiba (LeT). The goal was to understand whether there were any pure (or mixed) equilibria in which the group’s terrorist acts could be significantly reduced. The five players considered are: the US, India, the Pakistani military, the Pakistani civilian government, and the terrorist group LeT. Table I shows the actions the players were allowed to take.

When it comes to the application of game-theoretic reasoning to international strategic elements [Schelling 1960] with both state and non-state actors, the situation becomes much more complex because identifying the payoffs for different players is an enormous challenge and experts vary widely on what these payoffs are. To address

this application, we asked three internationally acknowledged world experts to give us payoff matrices—and we received three payoff matrices with substantial differences between them. Leveraging the theoretical constructs and results described above, we built the PREVE (Policy Recommendation Engine based on Vector Equilibria) software suite, and used it to identify approximate equilibria in the multiple payoff game induced by the three expert payoff matrices. We present key results produced by PREVE, and analyze their strengths and weaknesses from a policy perspective.

When it comes to the application of game-theoretic reasoning to international strategy [Schelling 1960] with both state and non-state actors, the situation becomes much more complex because identifying the payoffs for different players is an enormous challenge and experts vary widely on what these payoffs are. To address this application, we asked three internationally acknowledged world experts to give us payoff matrices—and we received three payoff matrices with substantial differences between them. Leveraging the theoretical constructs and results described above, we built the PREVE (Policy Recommendation Engine based on Vector Equilibria) software suite, and used it to identify approximate equilibria in the multiple payoff game induced by the three expert payoff matrices. We present key results produced by PREVE (Policy Recommendation Engine based on Vector Equilibria), and analyze their strengths and weaknesses from a policy perspective.

The paper begins with Section 2 introducing our LeT example briefly. As the LeT example is quite complex, a small toy example is also introduced. This toy example is used throughout the paper in order to illustrate the various definitions and algorithms we introduce. Section 3, our first formal section, consists of preliminaries which cover basic game-theoretic concepts. Section 5 formally defines our equilibrium concepts and presents bounds on computing them under various assumptions. Section 6 presents a QPTAS for the general case, when the number of players is constant. Building on this QPTAS, it gives efficient algorithms for computing such equilibria and experimentally validates them on simulated data. Section 7 gives a brief description of the computational system we built, called PREVE, and applies it to a real-world experimental five-player game used to model LeT. Section 7 also summarizes results from computing equilibria from three payoff matrices (created by area experts using open source data) and presents key policy results. Section 8 describes related work on game-theoretic models of terrorist group behavior as well as past policy recommendations on how the US and India should deal with LeT.

2. MOTIVATING EXAMPLES

In this section, we briefly describe the LeT application motivating this research. We also introduce a toy example that will be used throughout the paper to illustrate definitions, as the full LeT example can be too complex for that.

2.1. Reducing Terror Attacks by LeT

Lashkar-e-Taiba (LeT) is a terrorist group primarily funded by the Pakistani intelligence agency, the Inter-Services Intelligence [Winchell 2003]. Created in 1990, the group has carried out numerous terrorist attacks, the most spectacular of which was the November 2008 terrorist attack in Mumbai that targeted several sites including the iconic Taj Mahal hotel, killing 166 innocent civilians (as well as nine terrorists, while a tenth terrorist was captured). LeT has strong links to various other terrorist groups including Al-Qaeda, Indian Mujahideen, Jaish-e-Mohammed, Jabhat-al-Nusra in Syria, groups in Chechnya, Jemaa Islamiyah, as well as organized crime groups such as Dawood Ibrahim's D-Company. For instance, Al-Qaeda leader Khalid Sheikh

Mohammed was captured in an LeT safehouse in Pakistan. Given its technical sophistication and the support of a sophisticated intelligence agency, LeT is viewed as a major threat by both the US and India—both in terms of operations they might carry out themselves and in terms of training and logistics support they might provide to other groups that carry out such attacks.

In order to reduce terror attacks by LeT, we developed a five-player game. The players considered are the United States (US), India, Pakistan’s government, Pakistan’s military, and Lashkar-e-Taiba (LeT). We recall that Table I, presented earlier, gives actions each player can take, and that—in addition to the actions below—each player can take the action none, which corresponds to doing nothing. We describe the other actions in depth here.

US Actions. The US can take three actions (and none).

- (1) The first is covert action against LeT. While we do not suggest specific operations, this action could be implemented in many ways including covert actions to undermine LeT’s leaders or covert actions to target LeT training camps. It is clear that the US is capable of such covert action as evidenced by recent events involving a CIA contractor called Raymond Davis who was arrested by the Pakistanis after a shootout in Lahore.
- (2) The US could also cut military and/or development support currently being given to Pakistan. According to the Congressional Research Service, the US provided \$1.727 billion in economic aid to Pakistan in FY2010.² In 2012, the US asked Congress for permission to ship almost \$3 billion to Pakistan with over half being military aid.³ Cutting some of this aid is an option the US has long considered, especially in view of US Admiral Mike Mullen’s assertions in 2011 about Pakistan’s ISI controlling the Haqqani terrorist network which in turn attacked the US embassy in Kabul.⁴
- (3) The US could also expand financial support for Pakistan. Pakistan’s educational system and economy are both in shambles and some have argued that additional development assistance would wean young people away from radical elements.

India’s Actions. As with the US, we study three actions (and none) that India might take. Similarly, there are many ways in which India could tactically implement these actions.

- (1) Like the US, India can also take different forms of covert action against LeT using methods similar to those listed above for the US.
- (2) India can also use coercive diplomacy in which diplomatic moves are used to coerce Pakistan. For instance, a credible threat can be used to warn Pakistan of the consequences of carrying out certain actions. For coercive diplomacy to be effective, the threat must be made publicly and must be credible [Schelling 1960]. Credible threats could include withholding water by diverting the headwaters of the Indus or by troop movements or simply by ramping up military spending which would place pressure on other parts of the Pakistani economy.
- (3) A third option we consider is one where India proposes some kind of peace initiative to Pakistan, e.g., granting some additional rights for back and forth movement between India and Pakistan, unifying families in Kashmir who were split up by the partition of Kashmir, and so forth.

²See “Pakistan-U.S. Relations: A Summary,” by K. Alan Kronstadt of the Congressional Research Service, May 16, 2011.

³<http://www.foxnews.com/topics/us-aid-to-pakistan-fy2012-request.htm>

⁴<http://www.nytimes.com/2011/09/23/world/asia/mullen-asserts-pakistani-role-in-attack-on-us-embassy.html?pagewanted=all&r=0>

Pakistan Military Actions. We study three possible actions for the Pakistani military, all related to their support for LeT.

- (1) The Pakistani military could implement a crackdown on LeT by arresting LeT members and/or closing down LeT's training camps, shutting down the logistical support for LeT operations in Jammu and Kashmir, and taking steps to interdict LeT-allied organizations like Jamaat-ud-Dawa. Pakistani security has, at times, cracked down on LeT, e.g., after the December 2001 parliament attack and the November 2008 attacks in Mumbai.
- (2) The Pakistani military could cut support to LeT by, e.g., arresting military officers who are illicitly supporting LeT and stopping military training of LeT personnel.
- (3) The Pakistani military could also expand support for LeT, e.g., by increasing its logistical and materiel support as well as financial support.

Pakistan Government Actions. We consider just two possible actions (in addition to none) by the civilian side of the Pakistani government (excluding the military side).

- (1) The Pakistani government could prosecute and arrest LeT personnel, as they have done periodically (though the leaders are usually released shortly thereafter).
- (2) The Pakistani government could choose to endorse LeT's social services program by routing government services through them. LeT runs many social services in Pakistan ranging from ambulances to hospitals, schools, and disaster relief programs.

Lashkar-e-Taiba's Actions. In the case of LeT, we considered three actions (in addition to the none action).

- (1) LeT could launch a major attack. We already know from the November 2008 Mumbai siege that they have the capability and logistical support to execute such attacks.
- (2) LeT could hold attacks (but not major ones), similar to those periodically carried out by them in Kashmir where military and civilian personnel are frequently targeted.
- (3) LeT could do something dramatic like eliminate its armed wing, give up its weapons, and publicly renounce violence. Though extremely unlikely, this is still worth listing as a possible action.

2.2. A Toy Example

We now introduce a small example that will be used to illustrate formal concepts and definitions as they are introduced later in the paper. Consider a very simple game consisting of two players, a terrorist group T and a government G . Suppose the terrorist group can carry out two actions (terror-attack and peace) and the government can carry out two actions (CT-ops and peace). Here, CT-ops denotes some traditional counter-terror operations such as killing and arresting group members. Experts are divided on the values of these actions to each player and thus provide two payoff matrices, PM_1 and PM_2 .

	CT-ops	peace		CT-ops	peace
terror-attack	(-5, -5)	(3, -10)	terror-attack	(-5, 6)	(3, 3)
peace	(-8, 6)	(0, 0)	peace	(-5, -5)	(-5, 2)
	PM_1			PM_2	

Much of our analysis is for games with payoffs in $[0, 1]$. Note that a scaled and translated version of the above matrices that does not alter equilibria of the game can be constructed. The modified payoff matrices (rounded to hundredths) are given below.

	CT-ops	peace		CT-ops	peace
terror-attack	(0.31, 0.31)	(0.81, 0)	terror-attack	(0, 1)	(0.73, 0.73)
peace	(0.12, 1)	(0.62, 0.62)	peace	(0, 0)	(0, 0.64)
Scaled PM_1			Scaled PM_2		

In each of these tables, the rows show terror group T 's actions and the columns show the government G 's actions. For example, the entry $(3, -10)$ in PM_1 says that the payoff to the terror group is 3 and the payoff to the government is -10 when the terror group performs terror-attack and the government proposes peace.

We will use this simple motivating example to illustrate various concepts in this paper.

3. TECHNICAL PRELIMINARIES: APPROXIMATE EQUILIBRIA

In this section, we first review common game-theoretic models and equilibrium concepts (§3.1), then build on them to define approximate equilibria in games with multiple payoff functions (§4).

3.1. Approximate Equilibria in Games with a Single Payoff Function

We consider simultaneous multiplayer games. Let $[n] = \{1, 2, \dots, n\}$ be the set of players and $[m] = \{1, 2, \dots, m\}$ be the set of actions for each player. Let Δ_m be the simplex $\{(x_1, x_2, \dots, x_m) \mid \sum_{i \in [m]} x_i = 1, x_i \geq 0, \forall i \in [m]\}$.⁵

For any player j , any $\sigma^j \in \Delta_m$ is a probability distribution over the set of actions $[m]$; thus, σ^j is called a *strategy* for player j . If $\sigma^j = (x_1, x_2, \dots, x_m)$, then x_i is the probability that player j will perform action i . When all but one of the x_i 's in σ^j are 0, σ^j is called a *pure strategy*; otherwise, it is called a *mixed strategy*. In mixed strategies, a player probabilistically chooses which action to take. Note that we will calculate these mixed strategies from the multiple payoff matrices provided by experts. They are not inputs to our algorithms (and so experts do not have to provide them); they are outputs generated by our system.

We use Δ to denote the set $\prod_{j=1}^n \Delta_m$. Any $\sigma \in \Delta$ is called a *strategy profile* for a game. If $\sigma = (\sigma^1, \dots, \sigma^n) \in \Delta$, then σ^j denotes the strategy of the player j . For convenience, we can represent a strategy profile σ as (σ^j, σ^{-j}) , where σ^j represents the strategy of player j and σ^{-j} represents strategies for the rest of the players.

Example 3.1. Consider the toy example given in Section 2.2. An example *pure strategy* for the government G is to play action CT-ops. Similarly, a pure strategy for the terror group T is to play action terror-attack. An example of a *mixed strategy* for G is to play action CT-ops and action peace with probabilities of $\frac{1}{3}$ and $\frac{2}{3}$, respectively. Similarly, T could play action terror-attack and action peace with probabilities $\frac{1}{2}$ and $\frac{1}{2}$, respectively. The above mixed strategies for G and T together form a *strategy profile* for the game.

The *payoff* for a player j is a function $u_j : \Delta \mapsto [0, 1]$. In this section, we assume (without loss of generality) that all payoffs are in the unit interval $[0, 1]$. We now define a basic building block of game theory, the Nash equilibrium.

Definition 3.2. A strategy profile σ is a *Nash equilibrium* iff:

$$u_j(\sigma^{j'}, \sigma^{-j}) \leq u_j(\sigma) \quad \forall \sigma^{j'} \in \Delta_m, j \in [n]$$

⁵For ease of exposition, we assume identical action spaces $[m]$ for each player $i \in [n]$. We note that the model can easily be adapted to handle non-identical action spaces; indeed, our experimental results (Section 7) are performed in such a heterogeneous model.

Thus, a strategy profile is a Nash equilibrium if no player has incentive to deviate from his strategy, assuming all other players play their respective strategies. Classical game theory assumes that players are rational. Hence, players can reason about one another and identify the Nash equilibria that are possible and then typically play actions consistent with one such Nash equilibrium. As Schelling [1960] observes, a good amount of work may also be invested by players in “prepping” the game so that certain strategy profiles are excluded from being equilibria.

Example 3.3. Consider the mixed strategy given in Example 3.1. G plays action CT-ops and action peace with probabilities $\frac{1}{3}$ and $\frac{2}{3}$, respectively. Similarly, T plays action terror-attack and action peace with probabilities $\frac{1}{2}$ and $\frac{1}{2}$, respectively. In this case, as per the payoff matrix PM_1 (defined in Section 2.2), the payoff for G is $-5 * \frac{1}{3} * \frac{1}{2} + -10 * \frac{2}{3} * \frac{1}{2} + 6 * \frac{1}{3} * \frac{1}{2} + 0 * \frac{2}{3} * \frac{1}{2} = -\frac{19}{6}$. We note that the payoff for a given strategy profile is the expected payoff given players draw actions independently at random according to their respective strategies. A Nash equilibrium for the same game is for G to play CT-ops with probability 1 and for T to play terror-attack with probability 1, resulting in a payoff of -5 for both players.

Since Nash equilibria are notoriously difficult to compute [Chen and Deng 2006; Daskalakis et al. 2006a], recent work has focused on finding *approximate* Nash equilibria. We use a well-known notion of an approximate Nash equilibria.

Definition 3.4. A strategy profile σ is an ϵ -approximate Nash equilibrium for some $0 \leq \epsilon \leq 1$ iff:

$$u_j(\sigma^{j'}, \sigma^{-j}) \leq u_j(\sigma) + \epsilon, \forall \sigma^{j'} \in \Delta_m, j \in [n]. \quad (1)$$

A stricter notion of an approximate Nash equilibrium is the well-supported approximate Nash equilibrium. Let $S(\sigma)$, the *support* of a strategy $\sigma \in \Delta_m$, be the set $S(\sigma) = \{i \mid \sigma_i > 0\}$. Intuitively, the support of σ is the set of actions that are executed with nonzero probability. Daskalakis, Goldberg, and Papadimitriou [2006a] define a well-supported approximate Nash equilibrium as follows.

Definition 3.5. Suppose $0 \leq \epsilon \leq 1$ is a real number. A strategy profile σ is a well-supported ϵ -approximate Nash equilibrium iff:

$$u_j(e_i, \sigma^{-j}) \leq u_j(e_l, \sigma^{-j}) + \epsilon \quad \forall \sigma^{j'} \in \Delta_m, i \in [m], \\ l \in S(\sigma^j), j \in [n]$$

In other words, for a strategy to be a well-supported ϵ -approximate Nash equilibrium, every player’s incentive to deviate from his equilibrium strategy is very small (less than a utility of ϵ). Put another way, this stronger constraint is imposed by forcing any pure strategy $i \in [m]$ to be given 0 probability (i.e., to not be included in the support $S(\sigma)$) if its expected payoff is more than ϵ below the best response payoff.

Definition 3.4 and Definition 3.5 both define approximate Nash equilibria that are additive in nature (due to the $+\epsilon$ term in the right side of the definition). A multiplicative (relative) approximation can be defined as follows, due to [Daskalakis et al. 2006a].

Definition 3.6. A strategy profile σ is a *well-supported relative ϵ -approximate Nash equilibrium* for $0 \leq \epsilon \leq 1$ iff $\forall j \in [n], i \in [m]$:

$$(1 - \epsilon)u_j(e_i, \sigma^{-j}) \leq u_j(e_l, \sigma^{-j}), \forall l \in S(\sigma^j) \quad (2)$$

The following example illustrates these different notions of approximate equilibria.

Example 3.7. Consider the strategy profile from Example 3.1. G plays action CT-ops and action peace with probabilities $\frac{1}{3}$ and $\frac{2}{3}$, respectively. Similarly, T plays action terror-attack and action peace with probabilities $\frac{1}{2}$ and $\frac{1}{2}$, respectively. For T , this is an ϵ -approximate Nash equilibrium with $\epsilon = 1.5$ since the expected payoff for T is -1.17 and deviation to action terror-attack leads to a payoff of 0.33 . This is a well-supported ϵ -approximate Nash equilibrium for T with $\epsilon = 3$ because the payoff for action peace, which is in the support of T 's strategy, is -2.67 and deviation to terror-attack leads to a payoff of 0.33 —a gain of 3.00 .

4. APPROXIMATE EQUILIBRIA IN GAMES WITH MULTIPLE PAYOFF FUNCTIONS

In this section, we merge together the ideas of (well-supported) approximate Nash equilibria and Shapley's vector payoffs in an effort to combine multiple conflicting experts' knowledge of payoffs.

Definition 4.1. A simultaneous game with multiple payoff functions (SGM) is a triple $G = (n, m, U)$ where $[n]$ is a set of players, $[m]$ is the set of actions for each player in $[n]$, and $U = (U_1, U_2, \dots, U_f)$ consists of f ordered sets of payoff functions $U_k = (u_1^k, u_2^k, \dots, u_n^k), \forall k \in [f]$.

Intuitively, an SGM G can be viewed as f different games specified over a set of players, over the same strategy space, with payoff functions for players given by $U_k, k \in [f]$. We refer to these f individual simultaneous games as *constituent games* of G . Throughout this paper, we use the variable f to denote the number of payoff matrices considered—which is also equal to the number of constituent games in a SGM or a ZSGM (a zero-sum version of an SGM defined later in Section 5.1).

For instance, in our toy example, the game G consists of two different constituent games, one corresponding to each of the two payoff matrices.

We now merge the idea of an approximate Nash equilibrium (Definition 3.4) with that of Shapley's vector payoffs.

Definition 4.2. A strategy profile σ is a *multiple ϵ -approximate Nash equilibrium* of an SGM (n, m, U) , iff it is an ϵ -approximate Nash equilibrium for each of its constituent games. Specifically, for all $k \in [f]$:

$$u_j^k(\sigma^{j'}, \sigma^{-j}) \leq u_j^k(\sigma) + \epsilon, \forall \sigma^{j'} \in \Delta_m, j \in [n] \quad (3)$$

Building on Definition 4.2, we also combine well-supported approximate Nash equilibria (Definition 3.5) with vector payoffs.

Definition 4.3. A strategy profile σ is a *well-supported multiple ϵ -approximate Nash equilibrium* of an SGM (n, m, U) , iff it is a well-supported ϵ -approximate Nash equilibrium for each of its constituent games. That is, for all $k \in [f]$:

$$u_j^k(e_i, \sigma^{-j}) \leq u_j^k(e_l, \sigma^{-j}) + \epsilon \quad \forall \sigma^{j'} \in \Delta_m, i \in [m], \\ l \in S(\sigma^j), j \in [n]$$

Finally, we can define the multiplicative version of Definition 4.3 as well.

Definition 4.4. A strategy profile σ is a *well-supported multiple relative ϵ -approximate Nash equilibrium* of an SGM (n, m, U) iff it is a well-supported relative ϵ -approximate Nash equilibrium for each of its constituent games. Thus, for all $k \in [f], j \in [n], i \in [m]$:

$$(1 - \epsilon)u_j^k(e_i, \sigma^{-j}) \leq u_j^k(e_l, \sigma^{-j}), \forall l \in S(\sigma^j) \quad (4)$$

We now provide an example of a well-supported multiple ϵ -approximate Nash equilibrium in the context of our toy game.

Example 4.5. For the game defined in Section 2.2, consider the strategy profile where T plays action `terror-attack` with probability 1 and G plays action `CT-ops` with probability 1. For both the payoff matrices, PM_1 and PM_2 , this is a Nash equilibrium. Therefore, the given strategy profile is a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = 0$. Consider another strategy profile, where G plays action `CT-ops` and action `peace` with probabilities $\frac{1}{3}$ and $\frac{2}{3}$, respectively. Similarly, T plays action `terror-attack` and action `peace` with probabilities $\frac{1}{2}$ and $\frac{1}{2}$, respectively. This is a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = 5.5$. This is because in one of the payoff matrices, PM_1 , the payoff G receives from action `peace` in support is -5 and deviating to action `CT-ops` leads to a payoff of 0.5 leading to a gain of $\epsilon = 5.5$.

A well-supported multiple ϵ -approximate Nash equilibrium is “close” in payoff for each player to a (Nash or approximate Nash) equilibrium in the constituent game corresponding to each payoff matrix in the SGM. A well-supported multiple ϵ -approximate Nash equilibrium closely approximates equilibrium situations irrespective of which of the several experts’ payoff matrices is used—it is a robust.

For notational convenience, in the experimental section of this paper, we will refer to well-supported multiple ϵ -approximate Nash equilibria computed using only $U' \subseteq U$ payoff functions as (ϵ, ℓ) -*equilibria*, where $|U'| = \ell$. Such equilibria computed with the full set U are simply written as ϵ -*equilibria*.

5. APPROXIMATE EQUILIBRIA IN SIMULTANEOUS GAMES WITH MULTIPLE PAYOFF FUNCTIONS

We begin by analyzing two fairly constrained cases, zero-sum games (§5.1) and rank 1 games (§5.2). We then relax these assumptions, providing results for low-rank games (§5.3), which will later lead into results on general games where the number of players is constant (§6.2).

5.1. Zero-sum Games with Multiple Payoffs

We begin by extending the well-known linear program (LP) for computation of an exact Nash equilibrium in zero-sum games to the computation of an *approximate* Nash equilibrium, and subsequently use it to design an algorithm to compute multiple payoff equilibria in such games. We will focus on the zero-sum equivalent of a simultaneous game with multiple payoff functions, as defined by Definition 5.1.

Definition 5.1. A *zero-sum simultaneous game with multiple payoff functions (for two players)* (ZSGM) is an SGM $(2, m, U)$, with ordered set of payoff functions $U = (u^1, u^2, \dots, u^f)$ such that u^k ($-u^k$) is the payoff function for player 1 (player 2) $\forall k \in [f]$. For convenience, we denote such games $G = (m, U)$.

Note that ZSGMs are limited to just two players.

Let (m, U) , where $U = (u^1, u^2, \dots, u^f)$, be a ZSGM. Let $P = (r_1, r_2, \dots, r_f)$, $r_i \in [0, 1]$, $\forall i \in [f]$. Consider the following LP:

$$LP^f(U, P, \epsilon)$$

$$\begin{aligned}
& \sum_{i \in [m]} \sigma_i^1 = 1 \\
& \sigma_i^1 \geq 0, \forall i \in [m] \\
& \sum_{i \in [m]} \sigma_i^2 = 1 \\
& \sigma_i^2 \geq 0, \forall i \in [m] \\
& \sum_{i \in [m]} \sigma_i^1 u^k(e_i, e_j) \geq r_k - \epsilon, \forall j \in [m], k \in [f] \quad (5) \\
& \sum_{i \in [m]} \sigma_i^2 u^k(e_i, e_j) \leq r_k + \epsilon, \forall j \in [m], k \in [f] \quad (6)
\end{aligned}$$

Here, the first four equations are required because σ^1 and σ^2 are distributions over actions of players 1 and 2 respectively. Equations (5) and (6) are required because we want players to play strategies that are approximate best responses to each of the constituent games of the ZSGM.

For $f = 1$, this LP applies to a zero-sum game with scalar payoff function $u = u^1$. For this special case, we call this LP, $LP^1(u, r, \epsilon)$. Lemma 5.2 and Lemma 5.3 show that for the single payoff function case, the linear program LP^1 computes approximate Nash equilibria for zero-sum games. Thus, our framework neatly extends approximate Nash equilibria to the case when there are vector-valued payoffs. The next result states that any solution to the linear program given above yields an approximate Nash equilibrium.

LEMMA 5.2. *Any feasible solution to $LP^1(u, r, \frac{\epsilon}{2})$ is an ϵ -approximate Nash Equilibrium for a zero-sum game with u as payoff function for player 1.*

The following result says that every ϵ -approximate Nash equilibrium is a solution of the linear program LP given above.

LEMMA 5.3. *Any ϵ -approximate Nash equilibrium strategy profile (σ^1, σ^2) for a zero-sum game with payoff function u for player 1 such that payoff for player 1 is in $[r - \epsilon, r + \epsilon]$, $\epsilon \geq 0$ is a feasible solution to $LP^1(u, r, 2\epsilon)$.*

Lemma 5.4 and Lemma 5.5 below extend the above results (which apply when $f = 1$, i.e., when there is only one payoff function) to the case of zero-sum games with multiple payoff functions. The first result, analogous to Lemma 5.2, states that solutions of the above LP are multiple payoff ϵ -approximate equilibria.

LEMMA 5.4. *Any feasible solution to $LP^f(U, P, \frac{\epsilon}{2})$ is a multiple payoff ϵ -approximate equilibrium for the ZSGM (m, U) .*

The next result, analogous to Lemma 5.3 states that for every multiple payoff ϵ -approximate equilibrium, there is a corresponding solution of the above LP.

LEMMA 5.5. *Let $\sigma = (\sigma^1, \sigma^2)$ be a strategy profile that is a multiple payoff ϵ -approximate equilibrium for the ZSGM (m, U) . Let $P = (r_1, r_2, \dots, r_f)$ be the vector of payoffs for player 1 for each of the constituent games of the ZSGM. Let $P' = (r'_1, r'_2, \dots, r'_f)$ be a vector such that $|r_i - r'_i| \leq \epsilon, \forall i \in [f]$. Then σ is a feasible solution to $LP^f(U, P, 2\epsilon)$.*

Thus, $LP^f(U, P, \epsilon)$ precisely captures the entire set of multiple payoff ϵ -approximate equilibria of our zero sum game.

Algorithm 1 presents a method to compute the set of all approximate ϵ -equilibria in the multiple payoff case. The algorithm uses an input e in order to regulate the approximation error factor, ϵ .

ALGORITHM 1: Approximate Multiple Payoff ϵ -Equilibrium in Zero-sum Games

Input: e , set of payoff functions $U = (u^1, u^2, \dots, u^f)$
Output: A set of LPs (see Theorem 5.6 for details)
 $S \leftarrow \{\frac{0}{e}, \frac{1}{e}, \dots, 1\}$
 $Payoffs \leftarrow \times_{i=1}^f S$ /* Cardinality: $(e + 1)^f$ */
 $LP_Set \leftarrow \emptyset$
for $P \in Payoffs$ **do**
 $LP_Set \leftarrow LP_Set \cup LP^f(U, P, \frac{1}{e})$
return LP_Set

The result below shows that Algorithm 1 computes certain types of multiple-payoff approximate equilibria.

THEOREM 5.6. *Algorithm 1 runs in time $O((e + 1)^f(2mf + 2m + 2))$ and outputs a set of LPs. Let S be the union of feasible regions of all LPs in the set returned by the algorithm. Then S satisfies the following conditions:*

- (1) All strategy profiles in S are approximate multiple payoff ϵ -equilibria with $\epsilon = \frac{2}{e}$.
- (2) All multiple payoff ϵ -equilibria with $\epsilon = \frac{1}{2e}$ are in S .

5.2. Multiplayer Games of Rank 1

We now deal with the problem of finding equilibria in low-rank multiplayer games with multiple payoffs. Our real-world LeT application is one example of a low rank multiplayer game.

The definition of rank that we use is equivalent to one given by Kalyanaraman and Umans [2007]. As is evident from recent papers (e.g., [Lipton et al. 2003; Kannan and Theobald 2007; Kalyanaraman and Umans 2007; Theobald 2009; Adsul et al. 2011]), games of low rank have generated considerable interest.

In this section, we first define *multiplayer games of rank K* . We then give a complete characterization of Nash equilibria for these games when $K = 1$ and use this characterization to compute *well-supported relative ϵ -approximate Nash equilibria*.

Definition 5.7. A *multiplayer game of rank K* is a game where the payoff function for each player is specified by K n -tuples of vectors, each of length m . Let $\alpha^{j,k} = (\alpha^{1,j,k}, \alpha^{2,j,k}, \dots, \alpha^{n,j,k})$ be the tuple specifying the payoff function for player j . Let $\rho = (e_{a_1}, e_{a_2}, \dots, e_{a_n})$ be a strategy profile with only pure strategies for each player, where a_i is the action for player i . Then, the payoff for player j is defined as:

$$u_j(e_{a_1}, e_{a_2}, \dots, e_{a_n}) = \sum_{k \in [K]} \prod_{i \in [n]} \alpha_{a_i}^{i,j,k} \quad (7)$$

For a strategy profile, $\sigma = (\sigma^1, \sigma^2, \dots, \sigma^n)$, payoffs are defined as usual. Let $A = \underbrace{[m] \times [m] \times \dots \times [m]}_{n \text{ times}}$ be the set of all possible combinations of actions of all players.

Then, payoffs are given by:

$$u_j(\sigma) = \sum_{a \in A} \prod_{i \in [n]} \sigma_{a_i}^i u_j(e_{a_1}, e_{a_2}, \dots, e_{a_n}) \quad (8)$$

$$= \sum_{a \in A} \prod_{i \in [n]} \sigma_{a_i}^i \sum_{k \in [K]} \prod_{l \in [n]} \alpha_{a_i}^{l,j,k} \quad (9)$$

We note that each payoff matrix is of rank at most K and it is input as a rank- K decomposition. As can be seen from Equation 7, the payoff matrix for player j is a sum of K terms and the k^{th} term is the tensor product of vectors in tuple $\alpha^{j,k}$. This is a complicated definition, so let us consider the case when $K = 2$. In this case, the payoff matrix is implicitly specified by two vectors, each of length m (the number of actions). Consider the strategy profile $\rho = (e_{a_1}, e_{a_2}, \dots, e_{a_n})$. This strategy profile tells us what actions each of the n players is taking. From Equation 9, in a rank $K = 2$ game, we compute the payoff for player 1 (i.e. $j = 1$) as

$$\prod_{i=1}^n \alpha_{a_i}^{i,1,1} + \prod_{i=1}^n \alpha_{a_i}^{i,1,2}.$$

The following example uses our toy example to illustrate rank K games.

Example 5.8. As an example, we give a rank-2 decomposition of payoff matrix PM_1 of player T given in Section 2.2. A rank-2 decomposition of the matrix is $\alpha^{1,1,1} = \{-1.76, -2.54\}$, $\alpha^{2,1,1} = \{3.05, -0.55\}$, $\alpha^{1,1,2} = \{-1.30, 0.90\}$, $\alpha^{2,1,2} = \{-0.28, -1.56\}$. It can be easily verified (up to rounding error) that for T , $PM_1 = \alpha^{1,1,1}(\alpha^{2,1,1})^T + \alpha^{1,1,2}(\alpha^{2,1,2})^T$. A similar decomposition of PM_1 for G is given by $\alpha^{1,2,1} = \{-3.25, 0.99\}$, $\alpha^{2,2,1} = \{1.92, 2.81\}$, $\alpha^{1,2,2} = \{0.67, 2.18\}$, $\alpha^{2,2,2} = \{1.88, -1.28\}$. Therefore, the game specified by PM_1 is a rank-2 game.

For the special case of rank 1 games, we drop the superscript k from vectors $\alpha^{i,j,k}$.

5.2.1. Nash Equilibria for Rank 1 Games. The following result presents a complete characterization of Nash equilibria for multiplayer games of rank 1.

LEMMA 5.9. *Let σ be a mixed strategy profile. Let $u'_{-j}(\sigma) = \prod_{i \in [n] \setminus \{j\}} (\sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j})$. Let the support of player j 's strategy be $S_j = \{l \mid \alpha_l^{j,j} = \max(\alpha^{j,j})\}$. σ is a Nash equilibrium iff:*

$$u'_{-j}(\sigma) > 0 \implies \text{support}(\sigma^j) \subseteq S_j$$

In the rest of this section, we assume that all players have a non-zero payoff at equilibrium, since, for any general multiplayer game, if a player has a zero payoff at equilibrium, which is the minimum possible payoff for the game, then her maximum and minimum possible payoffs are both zero for any choice of action and the player is free to take any action in her equilibrium strategy.

5.2.2. ϵ -Approximate Multiple Payoff Equilibria for Rank 1 Games. We now give a characterization of well-supported relative ϵ -approximate Nash equilibria. Here we solve the multiplicative approximation problem which is harder than the additive approximation for normal form games [Daskalakis 2011].

LEMMA 5.10. *A strategy profile $\sigma = \langle \sigma^1, \sigma^2, \dots, \sigma^n \rangle$ is a well-supported relative ϵ -approximate Nash equilibrium (with a non-zero payoff for all players) for a multiplayer game of rank 1 with payoffs as specified in Section 5.2.1 iff:*

$$\alpha_i^{j,j} \geq (1 - \epsilon)(\max \alpha^{j,j}), \forall j \in [n], i \in S(\sigma^j) \quad (10)$$

We now prove the main result for this section and give an algorithm for the computation of well-supported multiple relative ϵ -approximate Nash equilibria in multiplayer games of rank 1. The following theorem is the main result for this section.

THEOREM 5.11. *Consider an SGM of rank 1 with f different payoff functions for each player. For $t \in [f]$, let payoff function k for player j be specified by tuple $(\alpha^{1,j,k}, \alpha^{2,j,k}, \dots, \alpha^{n,j,k})$. Then, a strategy profile $\sigma = (\sigma^1, \sigma^2, \dots, \sigma^n)$ is a well-supported multiple relative ϵ -approximate Nash equilibrium, iff $\forall j \in [n]$:*

$$\alpha_i^{j,j,k} \geq (\max \alpha^{j,j,k})(1 - \epsilon), i \in S(\sigma^j), k \in [f] \quad (11)$$

Algorithm 2 leverages this result to compute well-supported multiple relative ϵ -approximate Nash equilibria for rank-1 games.

ALGORITHM 2: well-supported multiple relative ϵ -approximate Nash equilibria in rank-1 games

Input: ϵ , payoff vectors $\alpha^{i,j,t}$

Output: Allowed actions in supports of a feasible strategy profile

for $i \in [m], j \in [n]$ **do**

if $\forall t \in [f], \alpha_i^{j,j,t} \geq (\max \alpha^{j,j,t})(1 - \epsilon)$ **then**
 | Add i to actions in support of player j 's strategies.

if there exists any player with empty support set **then** /*If any support is empty the profile is infeasible*/

 | **return** NULL

else

 | **return** The support sets constructed in the for loop

We use a variant of our toy example to explain Algorithm 2.

Example 5.12. As an example, consider a rank-1 game where Player 1 is T and Player 2 is G . Let a rank-1 PM_1 for T be $\alpha^{1,1,1} = \{1, 0.5\}, \alpha^{2,1,1} = \{0.5, 1\}$. Let, a similar PM_1 for G be given by $\alpha^{1,2,1} = \{1, 0.25\}, \alpha^{2,2,1} = \{1, 0.25\}$. Let PM_2 for T be given by $\alpha^{1,1,2} = \{0.75, 1\}, \alpha^{2,1,2} = \{0.75, 1\}$. Let PM_2 for G be given by $\alpha^{1,2,2} = \{0.75, 1\}, \alpha^{2,2,2} = \{1, 0.75\}$. Let $\epsilon = 0.5$. Then for PM_1 , all mixed strategies for T are well-supported relative ϵ -approximate Nash equilibria. For G only action CT-ops can be in the support. For PM_2 , all mixed strategies for both the players are well-supported relative ϵ -approximate Nash equilibria. Therefore, a strategy profile where T plays some mixed strategy and G plays action CT-ops is a well-supported multiple relative ϵ -approximate Nash equilibrium with $\epsilon = 0.5$.

5.3. Multiple Payoff Games of Low Rank

In this section we consider the general case of *multiplayer games of low rank* (Definition 5.7). We prove that a class of strategies called “uniform strategies” (which we will define shortly) can be used to compute approximate Nash equilibria for these games when the number of actions is small. We then leverage this result to design an algorithm that computes the set of all multiple payoff equilibria for such games.

In this and the next section, we focus only on uniform strategies. Uniform strategies provide a tradeoff between simplicity and optimality that may be valuable to the end user. For example, in the LeT game we study later in the paper, a policy prescription like “India should take covert action against LeT with probability 0.0071” may not be very useful to the end user. A simpler policy prescription that is *almost* as good may be a much better option. We now define a uniform strategy profile.

Definition 5.13. A strategy profile $\sigma = (\sigma^1, \sigma^2, \dots, \sigma^n)$ is a t -uniform strategy profile if, $\forall i \in [n], \forall j \in [m]$:

$$\sigma_j^i \in \left\{ \frac{\ell}{t} \mid \ell \in \{0, 1, \dots, t\} \right\}$$

Intuitively, a t -uniform strategy discretizes the $[0, 1]$ real-valued interval into t segments and considers two probabilities within the same segment to effectively be the same. Thus, as t gets bigger, we get finer granularity. Thus, our selection of t controls the *granularity* of the probability distribution on actions in a strategy. A smaller t leads to a coarse-grained, simple strategy whereas a larger t allows a more fine grained strategy that may be closer to an optimal strategy.

5.3.1. Approximate Nash Equilibria in Games of Low Rank. In this subsection, we constructively prove that a uniform strategy profile can be used to approximate a Nash equilibrium for multiplayer games of low rank. First, we state the following lemmas to help with the main result.

LEMMA 5.14. *Let α be a vector of length m such that each element of α is in $[0, 1]$. Let σ be vector of length m . Let σ' be a vector such that $|\sigma_i - \sigma'_i| \leq \epsilon, \forall i \in [m]$. Then $|\alpha^T \sigma - \alpha^T \sigma'| \leq m\epsilon$.*

LEMMA 5.15. *Let x_1, \dots, x_n be n reals such that $0 \leq x_i \leq 1, \forall i \in [n]$. Let x'_1, \dots, x'_n be n reals such that $0 \leq x'_i \leq 1, |x_i - x'_i| \leq \epsilon, \forall i \in [n]$. Then $|\prod_{i \in [n]} x_i - \prod_{i \in [n]} x'_i| \leq n\epsilon$.*

The main technical result of this subsection is that if a strategy profile is a well-supported multiple ϵ -approximate Nash equilibrium, then there exists a t -uniform strategy profile that is also a well-supported multiple ϵ -approximate Nash equilibrium with a slightly higher value of ϵ . However, the simpler lemma below—dealing with the single payoff case—provides the basis for the more complex theorem to follow.

LEMMA 5.16. *Let the strategy profile $\sigma = (\sigma^1, \sigma^2, \dots, \sigma^n)$ be a well-supported ϵ -approximate Nash equilibrium for the given game of rank k . Then there exists a t -uniform strategy profile σ' that is a well-supported $\epsilon + \frac{2(n-1)mk}{t}$ -approximate Nash equilibrium.*

We now extend Lemma 5.16 to the multiple payoff case pertaining to well-supported multiple ϵ -approximate Nash equilibria.

THEOREM 5.17. *Let the strategy profile σ be a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau$ for the given SGM, all of whose constituent games are rank k games. Then, there exists a t -uniform strategy profile σ' that is a well-supported multiple ϵ -approximate Nash equilibrium, with $\epsilon = \tau + \frac{2(n-1)mk}{t}$.*

6. COMPUTING MULTIPLE PAYOFF APPROXIMATE EQUILIBRIA

Building on the theoretical results of the last section, we now provide an efficient algorithm for computing well-supported multiple ϵ -approximate Nash equilibria in games where the number of players is constant. First, we present a grid search algorithm for computing equilibria (§6.1), and show that is efficient through a quasi-polynomial time approximation scheme (QPTAS) (§6.2). This algorithm is validated on simulated data in Appendix A, and on real data in the next section (§7).

6.1. Algorithm for Computation of Equilibria

We leverage Theorem 5.17 to present Algorithm 3 which searches over the space of all uniform strategy profiles and outputs those that are well-supported multiple ϵ -

ALGORITHM 3: well-supported multiple ϵ -approximate Nash equilibrium for constant rank games

Input: t , payoff functions for the SGM
Output: t -uniform well-supported multiple ϵ -approximate Nash equilibrium strategy profiles
 $S \leftarrow$ Set of all possible t -uniform strategies
 $E \leftarrow \emptyset$
 $\Sigma \leftarrow \times_{i=1}^n S$ /* Cardinality: $O(m^{kn})$ */
for $l \in [f]$ **do**
 /* f is the number of constituent games of the SGM */
 $E_l \leftarrow \emptyset$
 for $\sigma \in \Sigma$ **do**
 $isEquilibrium \leftarrow \text{TRUE}$
 for $j \in [n]$ **do**
 if not $isEquilibrium$ **then**
 break
 $payoff \leftarrow u_j^l(\sigma)$
 for $i \in [m]$ **do**
 $payoff_i \leftarrow u_j^l(e_i, \sigma_{-j})$
 if $payoff_i - payoff > \epsilon$ **then**
 $isEquilibrium \leftarrow \text{FALSE}$
 break
 if $isEquilibrium$ **then**
 $E_l \leftarrow E_l \cup \{\sigma\}$
return $E_1 \cap E_2 \cap \dots \cap E_f$

approximate Nash equilibria. Input parameters to the algorithm are t and payoff functions for the constituent games. We assume that each payoff function is given as an oracle, which, given the strategy profiles, returns a vector with payoffs for all players. The output of the algorithm is the set of all t -uniform strategy profiles which are well-supported multiple ϵ -approximate Nash equilibria for the given SGM.

The algorithm first chooses a strategy profile and checks if it is an equilibrium (e.g., by solving the linear programs presented earlier in the paper). For each payoff function in the list of f payoff functions, it then iteratively looks at pairs of players, trying to set payoffs that are sufficiently close to each other in an attempt to find an equilibrium. It iteratively adds any valid solutions found to the solution and returns the solution at the end. Via Theorem 5.17, this coarse grid search is guaranteed to find a “reasonable” overall equilibrium (with respect to the parameter t).

The following example illustrates Algorithm 3 on our running toy example.

Example 6.1. We note that a well-supported multiple ϵ -approximate Nash equilibrium exists for the game with $\epsilon = 0$. This is the common equilibrium for both payoff matrices when T plays action terror-attack and G plays action CT-ops. Thus, to guarantee $\epsilon = 0.2$, we require $t = 40$ for this game. For illustrative purposes, to avoid enumeration of all strategies required for $\epsilon = 0.2$, we use the algorithm as follows. We enumerate all t -uniform strategies and report only those strategies that are well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = 0.2$.

We note that strategy profiles given in rows 1 through 4 of the above table are common Nash equilibria with $\epsilon = 0$. All the other 3-uniform profiles have $\epsilon > 0.2$. Thus strategy profiles in the first 4 rows are all the 3-uniform well-supported multiple ϵ -approximate Nash equilibria of the given game.

Though this algorithm can be expected to yield reasonable running times for games of any rank, the guarantees shown in Theorem 5.17 only apply to low rank games. This algorithm has the added advantage that we do not need to compute the tensor

	<i>T</i> strategy		<i>G</i> strategy		ϵ for <i>T</i>		ϵ for <i>G</i>		Max ϵ
	terror-attack	peace	CT-ops	peace	PM_1	PM_2	PM_1	PM_2	
1	1.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00
2	1.00	0.00	0.67	0.33	0.00	0.00	0.00	0.00	0.00
3	1.00	0.00	0.33	1.00	0.00	0.00	0.00	0.00	0.00
4	1.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00
5	0.67	0.33	1.00	0.00	0.19	0.33	0.00	0.03	0.33
6	0.67	0.33	0.67	0.33	0.19	0.33	0.24	0.03	0.33
7	0.67	0.33	0.33	1.00	0.25	0.33	0.73	0.03	0.73
8	0.67	0.33	0.00	1.00	0.19	0.33	0.73	0.03	0.73
9	0.33	1.00	1.00	0.00	0.19	0.48	0.00	0.55	0.55
10	0.33	1.00	0.67	0.33	0.19	0.48	0.24	0.55	0.55
11	0.33	1.00	0.33	1.00	0.25	0.48	0.73	0.55	0.73
12	0.33	1.00	0.00	1.00	0.19	0.48	0.73	0.55	0.73
13	0.00	1.00	1.00	0.00	0.19	0.38	0.00	0.00	0.38
14	0.00	1.00	0.67	0.33	0.19	0.38	0.24	0.00	0.38
15	0.00	1.00	0.33	1.00	0.25	0.38	0.73	0.00	0.73
16	0.00	1.00	0.00	1.00	0.19	0.38	0.73	0.00	0.73

decomposition of the game matrix. As we will show in Section 6.2, uniform strategies are expected to provide good results on general games too.

6.2. A General Approximation Lemma for SGMs

We prove the existence of a QPTAS for SGMs when the number of players is constant (which is the case in our example game as well as in most real world strategic games). For this we first state and prove an approximation lemma (an extension to SGMs of Althöfer’s Approximation Lemma [1994]). Our approximation lemma states that if a well-supported multiple ϵ -approximate Nash equilibrium exists for an n -player game, then there is a well-supported multiple ϵ -approximate Nash equilibrium for the game using a t -uniform strategy with a slightly larger ϵ .

We note that the original version of Althöfer’s Approximation Lemma applies only to two-player bimatrix games and its straightforward application leads to a QPTAS for computing well-supported ϵ -approximate Nash equilibria for two-player games. However, our extension of the lemma to multiplayer games with multiple payoffs is not straightforward. In the multiplayer setting, the variables we consider are mutually overlapping products of independent random variables. Hence, to apply Hoeffding’s bound [1963], we iteratively discretize strategies of players.

LEMMA 6.2. *Consider a game with players in set $[n]$ and each player with actions in set $[m]$. Let $\sigma = (\sigma^1, \dots, \sigma^n)$ be any well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau$ for the game. Then, for any $j \in [n]$, there exists a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau + \delta$ in which the strategy of player j is t -uniform for $t = \frac{2 \log(2fmn)}{\delta^2}$.*

We now prove the main result of this section, Theorem 6.3. This result is stronger than Theorem 5.17 for the general case (applied directly to the low rank case) due to its logarithmic dependence on m and f .

THEOREM 6.3. *Let the strategy profile σ be a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau$ for the given SGM. Then, there exists a t -uniform strategy profile that is a well-supported multiple ϵ -approximate Nash equilibrium, with $\epsilon = \tau + \delta$ where $t = \frac{2n^2 \log(2fm(n-1))}{\delta^2}$.*

Table II. Statistics on number of $(\epsilon, 2)$ equilibria found.

ℓ	ϵ	#Eq. found	#Eq. without LeT attacks
2	0	252	6
2	0.1	357	6
2	0.2	1696	9
2	0.3	13925	42

Thus, in Algorithm 3, if we set $t = \frac{2n^2 \log(2fm(n-1))}{\delta^2}$, we are sure to find a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau + \delta$ given that at least one well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau$ exists. The runtime is then $O(m^{nt})$. Thus, the same algorithm works for both general and low rank case, albeit with different performance guarantees.

7. POLICY ANALYSIS RESULTS

We developed the Policy Recommendation Engine based on Vector Equilibria (PREVE) using the equilibrium concepts and algorithms described earlier. Using PREVE, we were able to analyze the Lashkar-e-Taiba (LeT) application described in Section 2.

We first obtained payoff matrices from three experts in the politics of South Asia and LeT in particular; to avoid bias, none had any background in game theory and none had ethnic origins in the Indian subcontinent. Two were retired US State Department employees with over 30 years of knowledge of negotiations in the region. The third was the author of three well-known books on terrorism. The payoff matrices were created completely independently using *open source* information as well as expertise of these experts by following a set of instructions on what payoff values meant.

As described earlier, for notational convenience, we will refer to well-supported multiple ϵ -approximate Nash equilibria computed using only $U' \subseteq U$ payoff functions as (ϵ, ℓ) -*equilibria*, where $|U'| = \ell$. Such equilibria computed with the full set U are simply written as ϵ -*equilibria*.

Before presenting the policy implications of the results generated by PREVE, we present a summary of the (ϵ, ℓ) -equilibria we found in Table II. We limit the equilibria presented to those where LeT does not attack. *No such $(\epsilon, 3)$ -equilibria were found for $\epsilon \leq 0.5$, so we focus on the case when $\ell = 2$.* In the case of mixed equilibria, we list an equilibrium as having no LeT attacks when the probability of LeT attacking (action `attack`) or holding its current set of attacks (action `hold`) is 25% or less.

We found no $(0, 3)$ -equilibria where LeT did not perform violent actions, but we did find the following:

- (1) There were 20 $(0, 2)$ -equilibria in which experts #1 and #3 agreed, 218 $(0, 2)$ -equilibria with experts #2 and #3 agreeing, and 14 $(0, 2)$ -equilibria in which experts #1 and #2 agreed.
- (2) Of these 252 $(0, 2)$ -equilibria, there were just six in which LeT did not carry out attacks. There were no $(0, 2)$ -equilibria involving experts #1 and #2 in which LeT did not carry out attacks. Table III summarizes the actions present in these six situations. An equilibrium named $E_{\epsilon, j, j'}$ is used to denote an $(\epsilon, 2)$ equilibrium in which the two experts who “agree” are j and j' .

In all six $(0, 2)$ -equilibria listed above where LeT stands down, the US cuts aid (development and military) to Pakistan, and India either carries out covert action against LeT or engages in coercive diplomacy. Moreover, in most $(0, 2)$ -equilibria, the Pakistani military must crack down on LeT (though there is one case where they may expand support) and additionally, the Pakistani government must mostly prosecute LeT leaders (though there are two cases where they could do absolutely nothing). When we look at experts #2 and #3, we see that there are only two $(0, 2)$ -

Table III. All (ϵ, ℓ) equilibria with $\epsilon = 0, \ell = 2$ in which LeT does not attack.

Equil.	LeT	PakG	PakM	India	US
$E_{0,1,3}^1$	eaw	pros	crack	covert	cut
$E_{0,1,3}^2$	eaw	pros	crack	0.75: covert 0.25: coerce	cut
$E_{0,1,3}^3$	eaw	none	crack	coerce	cut
$E_{0,1,3}^4$	none	pros	support	covert	cut
$E_{0,2,3}^5$	eaw	pros	crack	coerce	cut
$E_{0,2,3}^6$	none	none	crack	covert	cut

Table IV. All (ϵ, ℓ) equilibria with $\epsilon = 0.1, \ell = 2$ which are not $(0, 2)$ -equilibria, and in which LeT does not attack.

Equil.	LeT	PakG	PakM	India	US
$E_{0.1,1,3}^7$	0.5: attack 0.5: none	pros	expand	covert	cut
$E_{0.1,1,3}^8$	0.25: attack 0.75: none	pros	expand	covert	cut
$E_{0.1,1,3}^9$	none	pros	expand	covert	cut

equilibria in which LeT does not attack—in one India takes covert action and the US cuts aid. In both scenarios, the Pakistani military cracks down on LeT—in one the Pakistani government prosecutes LeT personnel and does nothing in the other. When we do the same with experts #1 and #3, we see that there are four $(0, 2)$ -equilibria in which LeT does not attack. In all four, India takes either covert action or applies coercive diplomacy and the US cuts aid. In three cases, LeT eliminates its armed wing, while in another it does nothing. In the other two, LeT has a 50% (resp. 75%) chance of doing nothing and a 50% (resp. 25%) chance of attacking. In three cases, the Pakistani government prosecutes LeT personnel and does nothing in the fourth. In three of the cases, the Pakistani military cracks down on LeT, and in the one remaining case, it actually expands support for LeT. What these results may suggest is that *India should expand covert action against LeT with the US cutting financial aid to Pakistan at the same time* if the goal is to reduce violence by LeT.

- (3) We also looked at $(0.1, 2)$ -equilibria (i.e., where $\epsilon = 0.1$), which means that each player may lose up to 10% of their best utility while being near an equilibrium with 2 of the 3 experts in agreement. In this case, we see no $(0.1, 2)$ -equilibria involving experts #1 and #2 where LeT does not attack. But with experts #1 and #3, and experts #2 and #3, we do see such equilibria. As all $(0, 2)$ -equilibria continue to be $(0.1, 2)$ -equilibria, we only show new $(0.1, 2)$ -equilibria in Table IV.

With $\epsilon = 0.1$, we only get three new equilibria as compared to Table III. In all of these, the US needs to cut aid to Pakistan and India needs to carry out covert action against LeT. As in the previous table, this requires that the Pakistani government prosecute LeT. Even with an expansion in Pakistani military support for LeT, this provides hope that covert action on India's part and cuts in US aid to Pakistan will lead to reduced terrorist attacks by LeT.

We now consider $(\epsilon, 2)$ -equilibria, for $\epsilon \in \{0.0, 0.1, 0.2\}$. Though we computed $(\epsilon, 2)$ -equilibria for $\epsilon = \{0.3, 0.4, 0.5, \dots\}$, all of these equilibria involve players giving up 30% or more of their payoffs—something that we think is unlikely.

Of the 252 $(0, 2)$ -equilibria, there were five equilibria in which the US cut aid, India carried out either covert operations against LeT or coercive diplomacy against Pakistan, and the Pakistani military cracked down on LeT. In every one of these situations,

LeT either eliminated its armed wing or did nothing, and the Pakistani government either prosecuted LeT or did nothing. Moreover, there are 24 $(0, 2)$ -equilibria in which the US cuts aid and India carries out either covert action or coercive diplomacy—and in 5 of these 24 equilibria, LeT either eliminated its armed wing or did nothing. However, the situation is more complex. In our data, we noticed that one expert's payoffs were significantly different from those of the other two. In fact, there were vastly more equilibria between experts #2 and #3 than between experts #1 and #2 or between #1 and #3, suggesting expert #1 was a bit of an outlier. If we only consider experts #2 and #3, then the proportion of “good” equilibria where LeT stands down with the US cutting aid to Pakistan and India either engages in covert action or coercive diplomacy against Pakistan rises to 5 out of only 14. Of course, other inducements not considered in this study can be used to get the Pakistani military to crack down on LeT.

We continued the same analysis of the 357 $(0.1, 2)$ -equilibria. There were 23 equilibria where the US cut aid and India took covert action. Of these, 6 equilibria led to LeT either disbanding its armed wing or doing nothing—good outcomes for peace. If we ignored expert #1 (who continued to be an outlier when we considered $(0.1, 2)$ -equilibria), the number of “good” equilibria remained the same, with fewer (20) overall equilibria. Again, when the Pakistani military cracked down on LeT, there was a 100% chance of LeT either eliminating its armed wing or getting rid of terrorism altogether.

When we look at the 1696 $(0.2, 2)$ -equilibria, we see a similar pattern. We had a total of 51 $(0.2, 2)$ -equilibria, of which LeT cut attacks in 9. There were only 51 of these $(0.2, 2)$ -equilibria in which the US cut aid and India took either covert action or engaged in coercive diplomacy. However, we note that when the Pakistani military also cracks down on LeT (in addition to the US and Indian actions just described), the majority (8 out of 9) of the remaining equilibria involve LeT eliminating its attacks.

The derivation of a policy when faced with multiple experts' opinions that diverge is a complex challenge. In the real world, policies are debated extensively by policy makers. Computer programs are unlikely to generate national security policies automatically—we view the role of our paper and PREVE as providing valuable input to such policy makers by identifying equilibria for all players across the multiple games that result when a simultaneous game is being played. We suggest that players choose to take actions that are common across most equilibria. For instance, in our game involving Lashkar-e-Taiba, the US should definitely cut aid to Pakistan because all the well-supported multiple approximate equilibria agree on this. Even though not all equilibria agree on what India should do, they agree that India should either take covert action or practice coercive diplomacy in all cases; hence, this is what we recommend to policy makers. Additional challenges are raised by the presence of multiple equilibria, defined below.

What to do with Multiple Equilibria. An important suggestion made by Schelling [1960] when multiple equilibria exist is to act quickly in order to eliminate undesirable equilibria. For example, when there are n equilibria and a player likes only m of them ($m < n$), then he can, in some cases, take unilateral action that excludes the other equilibria from being equilibria. For instance, in our LeT game, the Indian government can explicitly state that under no circumstances would they entertain a peace deal with Pakistan. By giving up this flexibility in a very vocal way—and making it clear that they cannot back down on this by taking other action, otherwise they would lose future elections—they eliminate all equilibria in which they propose a peace deal. An interesting technical question that we do not address in the paper is how a player can unilaterally act so that he eliminates equilibria that he deems less preferable. Unilateral action is possible in the real world because there is always a time gap—players do not play exactly at the same time, as the formalization in this (and many game theory) papers requires.

Limitations of our analysis. As with any application of a formal game-theoretic model to a real-world situation, our framework and subsequent analysis are limited in a variety of ways. One potential issue is that the actions listed above need not be mutually exclusive in practice. For example, the United States could choose to conduct covert operations against LeT in the same time period as cutting international aid to Pakistan’s government (i.e., perform both actions `covert` and `cut` in unison). Our analysis assumes exclusivity over chosen actions (e.g., equilibrium $E_{0,1,3}^2$ states that India has a 75% probability of performing only `covert` and a 25% probability of performing only `coerce`, and never both together). One way to address this is to include as new actions all combinations of non-exclusive actions; for example, include a new action `covert+cut` where the US performs both covert action against LeT and cuts support to Pakistan’s government, but do *not* include a new action `cut+expand` under the assumption that the US cannot both cut and expand aid to Pakistan’s government (i.e., `cut` and `expand` are *independent* or mutually exclusive). This adds to the complexity of the model, both in terms of the human expert-created payoff matrix creation and the approximate equilibrium computation. One could possibly address the former complexity bump by adapting techniques from the budgeted preference elicitation (e.g., Boutilier, Sandholm, and Shields [2004] in auctions, Nguyen et al. [2014] in Stackelberg games) or active learning and matrix completion communities (e.g., Sutherland et al. [2013]), and the latter complexity bump by parallelizing Algorithm 3 in the naïve way at either of the first two for loops.

Our choice of base solution concept—(approximate) Nash equilibrium—also warrants discussion of *equilibrium selection*, a well-known problem with many such concepts. There can be and typically are multiple equilibria; indeed, as discussed in this section, there were multiple favorable equilibria for our LeT problem, even when $\epsilon = 0$. Each agent must then *select* an equilibrium to play, a decision that is not currently included in our model (but has been addressed in different models in the game theory literature through general concepts like the correlated equilibrium). Furthermore, while our model assumes uncertainty over payoff matrices, it also makes some possibly unrealistic assumptions about symmetric information awareness across agents. For example, we assume that all agents know they are playing the same *game*, which may not be trivial in practice (e.g., the US may not know the full set of actions that LeT or the Pakistani military can take). This additional uncertainty could lead to agents adopting strategies that are not contained in *any* of the equilibria we compute. We leave these problems as future research.

8. RELATED WORK

We group our survey of related work into three sections: first, the purely theoretical aspects of computational game theory; second, the application of (computational and traditional) game theory to counter-terrorism and modeling conflict; and third, dealing with the purely social science study of Lashkar-e-Taiba.

8.1. Computational Game Theory

Games where each player has multiple payoffs have been studied before under many names such as vector-valued games [Shapley 1959], multi-criteria games [Mallozzi et al. 2008], games with multiple payoffs [Zeleny 1975], and multiple objective games [Zhao 1991]. However, past work mainly focuses on multiple payoffs as a way to model the situation where each player is trying to optimize many non-tradable and non-monetizable criteria simultaneously. For such games, Pareto equilibria [Borm et al. 1988] and its variants [Mallozzi et al. 2008] have been the solutions of choice.

However, the situation we consider compares alternate realities subscribed to by each expert. Hence, we are not interested in Pareto optimality.

As explained earlier, our motivation for this work is to analyze a simultaneous game when experts disagree on payoffs for players. Different experts (with unknown accuracy of prior knowledge) provide payoff functions for each player—and we expect experts to differ on such payoff functions because of subjective judgment in such applications. Since our work requires computation of approximate Nash equilibria that are common to all given payoffs, the problem is related to enumeration of Nash equilibria in the multiplayer setting.

The computation of even a single Nash equilibrium for two-player games is a hard problem [Chen and Deng 2006; Daskalakis et al. 2006a]. Enumeration of all Nash equilibria for a multiplayer game is likely to be an even harder problem [Avis et al. 2010]. There are also some hardness results known for computation of approximate Nash equilibria. It has been proven that it is unlikely that an FPTAS exists for the problem of finding Nash equilibria in a two-player game [Chen and Deng 2006]. Multiplicative approximation of Nash equilibria is also PPAD-complete for a constant approximation factor—even for two-player games [Daskalakis 2011].

Recently, there has been considerable progress in computation of approximate Nash equilibria for two-player games. The best known approximation factor for a polynomial time algorithm is 0.3393 [Tsaknakis and Spirakis 2007]. However, most recent work focuses on computation of a single Nash equilibrium for two-player games.

Theobald [2009] studies enumeration of Nash equilibria for two-player games of rank 1; however, that algorithm is not known to run in polynomial time. Lipton, Markakis and Mehta [2003] give the first QPTAS for computation of Nash equilibria in two-player and multiplayer games. However, the exponent depends on the inverse square of the approximation factor and on the square of the number of players and hence the algorithm is not feasible in practice. In fact, it has been proven by Feder, Nazarzaded and Saberi [2007] that, as far as brute force search over uniform strategy profiles is concerned, the runtime for these algorithms is tight. However, the above two results do indicate that, in general, a uniform grid search over the strategy space is a good heuristic for finding approximate Nash equilibria.

A particularly pertinent paper is that of Kalyanaraman and Umans [2007], which defines constant rank multiplayer games and gives a PTAS for finding approximate Nash equilibria for such games. We prove a structural theorem and also give a polynomial time algorithm for computation of Nash equilibria and well-supported multiple ϵ -approximate Nash equilibria for the rank 1 case. We also prove that when players have a small number of strategies to choose from, an assumption which holds for many real-world games, then a uniform strategy does well in the constant rank case.

We also note that our problem is different from simultaneous Bayesian games where players themselves are uncertain of the payoffs. An alternative view of the problem may be to model our problem as a Bayesian game where players are uncertain about their payoffs. However, Bayesian games require knowledge about priors which are hard to estimate for the real-world problems we are trying to model. In any case, we study the situation where experts differ in their perception of payoffs and try to find the approximate equilibria on which they all agree, and where we have *no prior* over the accuracy of each expert. We do not try to model the case where players themselves are uncertain of their payoffs. Our definitions of equilibria for multiple payoffs are related to the definition of *minimax-regret equilibrium* given by Hyafil and Boutilier [2004]. While their work defines the equilibrium so as to minimize the maximum regret, we seek to find all equilibria with a given bound for maximum regret.

8.2. Game Theory and the Study of Conflict

Though there has been extensive work on the use of game theory for political analysis, almost none of it involves large multiplayer games, and almost none of it involves the use of formal computational methods. The use of game theory to study conflict was pioneered by Schelling [1960], who developed a social scientist's view of how two-player conflicts—including terrorism—could be studied via game theory. Later, Bueno de Mesquita [2010] recounts how he used two-person games to predict various actions including one of interest in this project, namely that current US President Obama would not be able to stop Pakistani-based terrorism. Both these and similar efforts focus on two-player games; in contrast, the theory of equilibria in multiplayer games with multiple payoff matrices was not described by either of them. Lastly, this paper uses the LeT game proposed by Dickerson et al. [2011]. In contrast to this work, which used only one payoff matrix corresponding to the views of a single expert, we use a multiple payoff matrix model in this paper for which the relevant game theory and the resulting implications for dealing with LeT had to be completely reconsidered.

Ozgul et al. [2007] have studied the problem of detecting terror cells in terror networks and proposed a variety of algorithms such as the GDM and OGD methods. Similarly, Lindelauf et al. [2009] have studied the structure of terrorist networks and how they need to maintain sufficient connectivity in order to communicate while simultaneously maintaining sufficient disconnectivity in order to stay hidden. They model this tension between communication and covertness via a game-theoretic model. This same intuition led to the concept of covertness centrality [Ovelgönne et al. 2012] in social networks where a statistical (rather than game-theoretic) method is used to predict covert vertices in a network.

Sandler and Enders [2004] use the ITERATE data set of terrorist events to discuss how economic methods including both game theory and time series analysis can be used to propose policies for counter-terrorism. In an earlier survey [Enders and Sandler 1995], the same authors specify how game theory might be used to model target selection by terrorists. Major [2002] uses a mix of game theory, search, and statistical methods to model terrorism risk. None of these works provide a formal game-theoretic model involving both multiple players and multiple payoff matrices.

8.3. Research on and Analysis of Lashkar-e-Taiba

On the social science side, Clark [2010] was the first to study LeT from a military perspective. He argues that LeT has grown beyond the control of Pakistan and the Directorate for Inter-Services Intelligence (ISI), and that it will continue to grow with help from fringe elements in the Pakistani military establishment. He argues that India can only insulate itself from LeT-backed attacks by diminishing the internal threat posed by the Indian Mujahideen, an Indian group closely affiliated with LeT.

Tankel [2011b] wrote a detailed analysis of LeT based on years of field work and multiple visits to Pakistan to interview both LeT operatives as well as members of Pakistan's ISI. He provides a wonderful insight into LeT's origins, ideology, and operational structure, but does not include a policy analytics section specifically saying how to deal with the menace posed by LeT. John's excellent volume [John 2011] on the same topic provides another in-depth study of LeT but does not propose policies on how the US and/or India can collectively help reduce LeT attacks.

Virtually all past work on counter-terrorism policy is qualitative (see work by Mannes [2013] for an overview). A group of experts gather around a table, hypothesize about the impacts of different possible policies, and then decide which one to use. It is only recently that quantitative methods for generating policies against terror groups have started playing a role. Data mining approaches have been used to study

the Pakistani terror group Lashkar-e-Taiba [Subrahmanian et al. 2012] with considerable impact in the strategic policy community in both the US and India, both of whom have attended talks on the results. Subrahmanian et al. [2012] performs a data mining study of LeT involving 770 variables that are analyzed via data mining algorithms to learn the conditions under which LeT executes various types of attacks. It goes on to consider the problem of shaping the behavior of LeT by using abductive inference models. A subsequent book carries out a similar analysis of the group Indian Mujahideen (IM) [Subrahmanian et al. 2013]. Both books were followed by multiple talks and articles in the press predicting attacks by LeT and/or IM, with most predictions being accurate. Another excellent recent book on Pakistan in general by Bruce Riedel [2012], a former top CIA official who advised the last five US presidents on relations with India and Pakistan, lays much of the blame for terrorism out of Pakistan (including LeT terrorism) squarely at the doorstep of the Pakistani intelligence agency but does not address LeT attacks in particular.

9. CONCLUSIONS

In this paper, we showed how to merge vector payoffs [Shapley 1959] and well-supported ϵ -approximate equilibria [Daskalakis et al. 2006a; Daskalakis et al. 2006b] so as to handle the problem of efficient computation of equilibria in multiplayer games where multiple experts provide different payoff matrices, each capturing their own perception of reality. We present efficient algorithms to find such equilibria—as well as a QPTAS—and experimental results showing they work. The work is motivated by a real-world game we have built to formulate policies against the terror group Lashkar-e-Taiba (LeT) which carried out the 2008 Mumbai attacks. We then presented PREVE, a set of algorithms based on multiplayer game theory that extends a game developed earlier [Dickerson et al. 2011] to the case where there are multiple payoff matrices that reflect differing opinions of different experts. As a consequence, the resulting equilibria are much more robust to variations than the equilibria developed in [Dickerson et al. 2011] that are very sensitive to minor changes in the payoff matrix.

Pakistan is widely recognized as being one of the biggest threats to global security today because of several factors: (i) its nuclear arsenal, (ii) the large milieu of violent terrorist and extremist groups in the area with close ties to Pakistani intelligence, (iii) tensions with India, and (iv) a collapsing economy. In this paper, we have focused primarily on Pakistan-India relations, which India views primarily through the lens of terrorist acts in India that are backed by the Pakistani military and are usually operationally executed by LeT and/or its allies, like the Indian Mujahideen.

The PREVE theory, framework, and code have been developed in order to help policy-makers with an interest in peace in South Asia determine the best ways for the parties involved to move forward in order to reduce the threat of Lashkar-e-Taiba. Though we applied PREVE only to LeT in this paper, the theory is general and can be applied to any set of actors with any set of actions as long as one or more payoff matrices are available. In this paper, area experts used *open source* data to create payoff matrices for our five-player game.

PREVE may be improved in many ways. For instance, as noted by Nobel laureate Tom Schelling and many others, it is well known that real-world conflicts are rarely zero-sum games [Schelling 1960]. We have used certain zero-sum assumptions in this paper in order to simplify an already very complex framework. However, many parts of this paper do not require the zero-sum assumption and so it has been removed in those parts. Nonetheless, there are interesting alternative models that could be used. For instance, we could have players with uncertainty over the *game* they are playing, and who play according to strategies that are robust across many possible games. In this case, the players in question may not be actually playing any of the well-supported

approximate equilibria in the true game. An improved treatment to handle this case needs to be developed.

From a public policy perspective, the results of this paper may support three ideas.

- (1) The US must cut aid to Pakistan. There are no equilibria where LeT behaves well where the US is providing aid to Pakistan. However, we do not have a recommendation for exactly how much this cut should be—only that cuts need to be made.
- (2) India must engage in additional covert action against LeT and its allies and/or coercive diplomacy towards Pakistan. By cutting aid, the US would intuitively increase political and economic pressure on the Pakistani establishment, leading to a potential loss of support for the Pakistani military leadership amongst the Pakistani people. By engaging in covert action, India would put operational constraints on LeT, making attacks harder by “taking the fight to them” as the US has done against Al-Qaeda. By taking steps towards coercive diplomacy, India would concurrently increase pressure on the Pakistani government and military, complementing the US aid cuts proposed.
- (3) The key policy element is getting the Pakistani military to crack down on LeT, in conjunction with US cuts on aid to Pakistan and covert action/coercive diplomacy by India. The key question is how to induce the Pakistani military to crack down on LeT. An examination of the deep social, political, economic, and jihadist links that the Pakistani military has could lead to better understanding of the pressures that might induce them to crack down on extremist elements, many of whom they currently support.

PREVE is a codebase, not an operational system. Top politicians and policymakers are busy and are often more interested in white papers addressing their problem than learning how to use software systems. In our case, PREVE has been used to generate these results and then generate a report interpreting the results for policymakers. The results of this study have been disclosed to top government officials in both the US and Indian government. There is significant interest in continuing these studies.

REFERENCES

- Bharat Adsul, Jugal Garg, Ruta Mehta, and Milind Sohoni. 2011. Rank-1 bimatrix games: a homeomorphism and a polynomial time algorithm. In *ACM Symposium on Theory of Computing (STOC)*. ACM, 195–204.
- Ingo Althöfer. 1994. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra Appl.* 199 (1994), 339–355.
- Jai Asundi, Octavian Carare, and Kutsal Dogan. 2012. Competitive implications of software open-sourcing. *Decision Support Systems* 54, 1 (2012), 153 – 163. DOI: <http://dx.doi.org/10.1016/j.dss.2012.05.001>
- David Avis, Gabriel D. Rosenberg, Rahul Savani, and Bernhard von Stengel. 2010. Enumeration of Nash equilibria for two-player games. *Economic theory* 42, 1 (2010), 9–37.
- Peter E.M. Borm, S.H. Tijs, and J. van den Aarssen. 1988. *Pareto equilibria in multiobjective games*. Technical Report. Tilburg University.
- Craig Boutilier, Tuomas Sandholm, and Rob Shields. 2004. Eliciting bid taker non-price preferences in (combinatorial) auctions. In *AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 4. 204–211.
- Xi Chen and Xiaotie Deng. 2006. Settling the complexity of two-player Nash equilibrium. In *IEEE Symposium on Foundations of Computer Science (FOCS)*.
- Ryan Clarke. 2010. *Lashkar-i-Taiba: The Fallacy of Subservient Proxies and the Future of Islamist Terrorism in India*. Strategic Studies Institute, US Army War College.
- Constantinos Daskalakis. 2011. On the complexity of approximating a Nash equilibrium. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1498–1517.
- Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. 2006a. The complexity of computing a Nash equilibrium. In *ACM Symposium on Theory of Computing (STOC)*. ACM, 71–78.
- Constantinos Daskalakis, Aranyak Mehta, and Christos H. Papadimitriou. 2006b. A note on approximate Nash equilibria. In *Workshop On Internet And Network Economics (WINE)*. Springer, 297–306.
- B. Bueno de Mesquita. 2010. *The Predictioneer's Game: Using the Logic of Brazen Self-Interest to See and Shape the Future*. Random House.

- John P. Dickerson, Aaron Mannes, and V.S. Subrahmanian. 2011. Dealing with Lashkar-e-Taiba: A Multi-Player Game-Theoretic Perspective. In *International Symposium on Open Source Intelligence and Web Mining (OSINT-WM)*. 354–359.
- John P. Dickerson, Anshul Sawant, Mohammad Hajiaghayi, and V.S. Subrahmanian. 2013. PREVE: A Policy Recommendation Engine based on Vector Equilibria Applied to Reducing LeT's Attacks. In *International Symposium on Foundations of Open Source Intelligence and Security Informatics (FOSINT-SI)*.
- Walter Enders and Todd Sandler. 1995. Terrorism: Theory and applications. *Handbook of defense economics* 1 (1995), 213–249.
- Tomas Feder, Hamid Nazerzadeh, and Amin Saberi. 2007. Approximating nash equilibria using small-support strategies. In *ACM Conference on Electronic Commerce (EC)*. ACM, 352–354.
- Wassily Hoeffding. 1963. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* 58, 301 (1963), 13–30.
- Nathanael Hyafil and Craig Boutilier. 2004. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Conference on Uncertainty in Artificial Intelligence (UAI)*. 268–277.
- Wilson John. 2011. *Caliphate's Soldiers: The Lashkar-e-Tayyeba's Long War*. Amaryllis Press and the Observer Research Foundation.
- Shankar Kalyanaraman and Christopher Umans. 2007. Algorithms for playing games with limited randomness. In *Electronic Colloquium on Computational Complexity*. Springer, 323–334.
- Ravi Kannan and Thorsten Theobald. 2007. Games of fixed rank: a hierarchy of bimatrix games. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 1124–1132.
- Ebrahim Karami and Savo Glisic. 2010. Optimization of Wireless Multi-source Multicast Ad-Hoc Networks Using Asymmetric Matrix Games. In *IEEE International Conference on Communications (ICC)*. 1–5.
- Sarit Kraus, Jonathan Wilkenfeld, and Gilad Zlotkin. 1995. Multiagent negotiation under time constraints. *Artificial intelligence* 75, 2 (1995), 297–345.
- Roy Lindelauf, Peter Borm, and Herbert Hamers. 2009. The influence of secrecy on the communication structure of covert networks. *Social Networks* 31, 2 (2009), 126–137.
- Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. 2003. Playing large games using simple strategies. In *ACM Conference on Electronic Commerce (EC)*. ACM, 36–41.
- John A. Major. 2002. Advanced techniques for modeling terrorism risk. *Journal of Risk Finance, The* 4, 1 (2002), 15–24.
- Lina Mallozzi, Lucia Pusillo, and Stef Tijs. 2008. Approximate equilibria for Bayesian games. *J. Math. Anal. Appl.* 343, 2 (2008), 1098–1102.
- Aaron Mannes. 2013. Qualitative Analysis & Computational Techniques for the Counter-Terror Analyst. *Handbook of Computational Approaches to Counterterrorism* (2013), 83–98.
- Thanh H Nguyen, Amulya Yadav, Bo An, Milind Tambe, and Craig Boutilier. 2014. Regret-based Optimization and Preference Elicitation for Stackelberg Security Games with Uncertainty. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Michael Ovelgönne, Chanhyun Kang, Anshul Sawant, and V. S. Subrahmanian. 2012. Covertness Centrality in Networks. In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. 863–870.
- Fatih Ozgul, Julian Bondy, and Hakan Aksoy. 2007. Mining for offender group detection and story of a police operation. In *Proceedings of the Sixth Australasian Conference on Data Mining and Analytics*. Australian Computer Society.
- Karen Pittel and Dirk TG Rübbelke. 2012. Transitions in the Negotiations on Climate Change: From Prisoner's Dilemma to Chicken and Beyond. *International Environmental Agreements: Politics, Law and Economics* 12, 1 (2012), 23–39.
- Bruce O. Riedel. 2012. *Deadly embrace: Pakistan, America, and the future of the global Jihad*. Brookings Institution Press.
- Todd Sandler and Walter Enders. 2004. An economic perspective on transnational terrorism. *European Journal of Political Economy* 20, 2 (2004), 301–316.
- Thomas C. Schelling. 1960. *The Strategy of Conflict*. Harvard University Press.
- Lloyd S. Shapley. 1959. Equilibrium points in games with vector payoffs. *Naval Research Logistics Quarterly* 6, 1 (1959), 57–61.
- VS Subrahmanian, Aaron Mannes, RK Raghavan, and Animesh Roul. 2013. *Indian Mujahideen: Computational Analysis and Public Policy*. Springer.
- V.S. Subrahmanian, Aaron Mannes, Amy Sliva, Jana Shakarian, and John P. Dickerson. 2012. *Computational Analysis of Terrorist Groups: Lashkar-e-Taiba*. Springer New York.
- Dougal J Sutherland, Barnabás Póczos, and Jeff Schneider. 2013. Active learning and search on low-rank matrices. In *International Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, 212–220.
- Stephen Tinkel. 2011a. *Lashkar-e-Taiba: Past Operations and Future Prospects*. New America Foundation.
- Stephen Tinkel. 2011b. *Storming the World Stage: The Story of Lashkar-e-Taiba*. Columbia University Press.
- Thorsten Theobald. 2009. Enumerating the Nash Equilibria of Rank-1 Games. In *CRM Proceedings and Lecture Notes*, Vol. 48.
- Haralambos Tsaknakis and Paul G. Spirakis. 2007. An Optimization Approach for Approximate Nash Equilibria. In *Workshop On Internet And Network Economics (WINE)*, Vol. 4858. Springer, 42.

- Sean P Winchell. 2003. Pakistan's ISI: The Invisible Government. *International Journal of Intelligence and Counterintelligence* 16, 3 (2003), 374–388.
- Robert Woodley, Michael Gosnell, and Warren Noll. 2008. Adding Culture to Command Decisions. In *Proc. 13th Intl. Command and Control Research and Technology Symposium (ICCRTS)*.
- Milan Zeleny. 1975. Games with multiple payoffs. *International Journal of game theory* 4, 4 (1975), 179–191.
- Jingang Zhao. 1991. The equilibria of a multiple objective game. *International Journal of Game Theory* 20, 2 (1991), 171–182.

Received —; revised —; accepted —

Online Appendix to: Automated Generation of Counter-Terrorism Policies using Multi-Expert Input

ANSHUL SAWANT, University of Maryland
JOHN P. DICKERSON, Carnegie Mellon University
MOHAMMAD T. HAJIAGHAYI, University of Maryland
V.S. SUBRAHMANYAN, University of Maryland

A. ADDITIONAL EXPERIMENTS

In this section, we present experimental results for Algorithm 3 on simulated data. First, we present results showing the algorithm's running time and output on generated games. Second, we explore the relationship between various traits of the game and the percentage of strategies that are equilibria. The framework was implemented in about 700 lines of C++, and the experiments were run on a 4-CPU, 4-core Intel Xeon 3.4GHz machine with 64GB of RAM running Ubuntu 12.04.

To test the scaling properties of Algorithm 3, we built a game generator and varied the number of experts (each giving one set of payoff matrices), players, and actions per player. We also varied the granularity factor t when generating t -uniform strategies.

Figure 1 shows the runtime of Algorithm 3 on generated data as both the number of players and number of actions increase, for varying granularity factors. As expected, increasing the number of players (while holding the number of actions constant) hurts runtime significantly more than increasing the number of actions (while holding the number of players constant). Similarly, increasing the granularity factor t (shown on the x-axis) exponentially increases the number of possible strategy profiles over which the algorithm must iterate, resulting in large runtime increases. Future research would increase the algorithm's equilibrium-generation capabilities to games with many players and many actions.

Figure 2 quantifies the relationship between the ϵ -approximation threshold and the percentage of strategy profiles that are well-supported multiple ϵ -approximate Nash equilibria. Intuitively, increasing the slack in the approximation factor ϵ yields a higher

© 2014 ACM 2157-6904/2014/09-ART1 \$15.00
DOI: <http://dx.doi.org/10.1145/0000000.0000000>

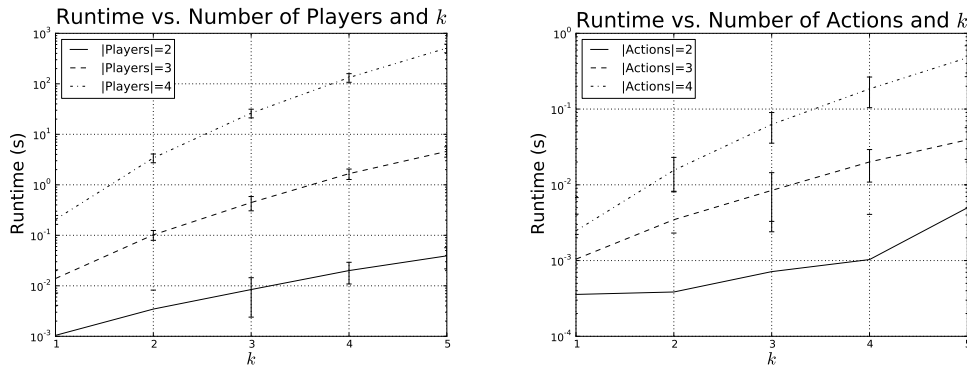


Fig. 1. Runtime as the number of players increases (left) and number of actions increases (right) for t -uniform factor $t \in \{1, \dots, 5\}$.

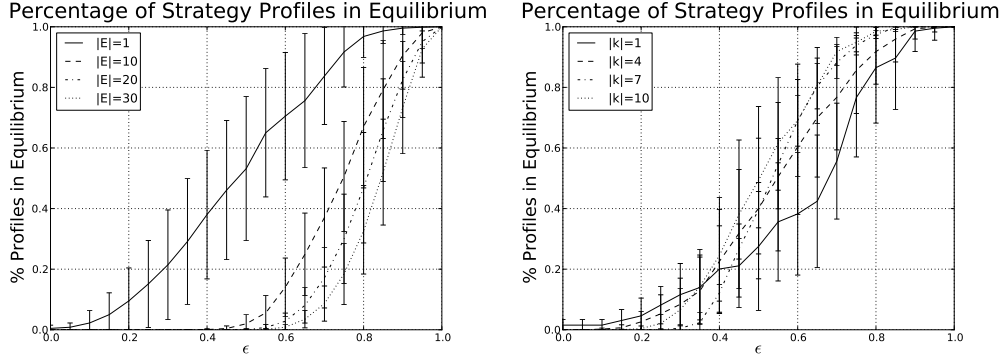


Fig. 2. Percentage of all sets of strategy profiles that are well-supported multiple ϵ -approximate Nash equilibria as the number of experts increases (left) and t -uniform factor increases (right), for $\epsilon \in \{0.0, 0.05, \dots, 1.0\}$.

percentage of strategy profiles being equilibria, while increasing the number of potential payoff matrices decreases this percentage of strategy profiles. The rate of increase of this line is highly dependent on the distribution of payoffs to each individual player. With random generation of payoffs, the increase is fairly steady; however, a more structured (e.g., real-world) payoff function would affect this trend. In Section 7, we considered such a real-world game.

B. PROOFS

In this section, we provide complete proofs for various theorems and lemmas in the main paper.

B.1. Proofs for Section 5.1

B.1.1. Lemma 5.2

PROOF. Let (σ^1, σ^2) be a feasible solution to the given LP. Let $p = \sum_{i \in [m]} \sum_{j \in [m]} \sigma_i^1 \sigma_j^2 u(e_i, e_j)$ be the payoff for player 1. The payoff for player 2 will be $-p$. Then, we have:

$$\begin{aligned}
 p &= \sum_{i \in [m]} \sum_{j \in [m]} \sigma_i^1 \sigma_j^2 u(e_i, e_j) = \sum_{i \in [m]} \sigma_i^1 \sum_{j \in [m]} \sigma_j^2 u(e_i, e_j) \\
 &\leq \sum_{i \in [m]} \sigma_i^1 (r + \epsilon) \quad (\text{from (6)}) \\
 &= r + \epsilon \because \sum_{i \in [m]} \sigma_i^1 = 1
 \end{aligned} \tag{12}$$

Similarly,

$$\begin{aligned}
 p &= \sum_{j \in [m]} \sigma_j^2 \sum_{i \in [m]} \sigma_i^1 u(e_i, e_j) \geq \sum_{j \in [m]} \sigma_j^2 (r - \epsilon) \quad (\text{from (5)}) \\
 &= r - \epsilon \because \sum_{j \in [m]} \sigma_j^2 = 1
 \end{aligned} \tag{13}$$

From (12) and (5):

$$\sum_{i \in [m]} \sigma_i^1(-u(e_i, e_j)) \leq -p + 2\epsilon, \forall j \in [m] \quad (14)$$

Similarly, from (13) and (6):

$$\sum_{j \in [m]} \sigma_j^2 u(e_i, e_j) \leq p + 2\epsilon, \forall i \in [m] \quad (15)$$

Since p and $-p$ are payoffs for given strategies and u and $-u$ are the payoff functions for players 1 and 2 respectively, the claim follows from (14), (15) and the definition of approximate Nash Equilibrium (Definition 3.4). \square

B.1.2. Lemma 5.3

PROOF. Let $p = \sum_{i \in [m]} \sum_{j \in [m]} \sigma_i^1 \sigma_j^2 u(e_i, e_j)$ be the payoff for player 1. Because its a zero-sum game, the payoff for player 2 will be $-p$. Then, from Definition 3.4 of approximate Nash Equilibrium:

$$\begin{aligned} & \sum_{i \in [m]} \sigma_i^1(-u(e_i, e_j)) \leq -p + \epsilon, \forall j \in [m] \\ \implies & \sum_{i \in [m]} \sigma_i^1 u(e_i, e_j) \geq p - \epsilon, \forall j \in [m] \quad (\text{multiplying by } -1) \\ \implies & \sum_{i \in [m]} \sigma_i^1 u(e_i, e_j) \geq r - \tau, \forall j \in [m] \quad \because p \geq r - \epsilon, 2\epsilon \leq \tau \end{aligned} \quad (16)$$

Similarly,

$$\begin{aligned} & \sum_{j \in [m]} \sigma_j^2 u(e_i, e_j) \leq p + \epsilon, \forall i \in [m] \\ \implies & \sum_{j \in [m]} \sigma_j^2 u(e_i, e_j) \leq r + \tau, \forall i \in [m] \quad \because p \geq r - \epsilon, 2\epsilon \leq \tau \end{aligned} \quad (17)$$

The other constraints in the LP are satisfied by any valid strategy profile. Thus, the claim follows from (16), (17) and the fact that (σ^1, σ^2) is a strategy profile. \square

B.1.3. Lemma 5.4

PROOF. Any feasible solution to $LP_MEAE(U, P, \frac{\epsilon}{2})$ is a feasible solution to $LP_EAE(u^i, r_i, \frac{\epsilon}{2}), \forall i \in [f]$ because constraints for $LP_MEAE(U, P, \frac{\epsilon}{2})$ are a superset of constraints for $LP_EAE(u^i, r_i, \frac{\epsilon}{2}), \forall i \in [f]$. Thus, from Lemma 5.2, the feasible solution is a strategy profile that is an ϵ -approximate Nash equilibrium for all constituent games of the ZSGM. The result then follows from Definition 4.2 of multiple ϵ -approximate Nash equilibrium. \square

B.1.4. Lemma 5.5

PROOF. From Lemma 5.3, any ϵ -approximate Nash equilibrium for zero-sum game with payoff matrix u^i such that payoff for player 1 is between $r_i - \epsilon$ and $r_i + \epsilon, \forall i \in [f]$ satisfies all constraints of $LP_EAE(u^i, r_i, 2\epsilon)$. Thus, the given equilibrium is feasible for $LP_EAE(u^i, r_i, 2\epsilon), \forall i \in [f]$. Hence the given equilibrium satisfies $LP_MEAE(U, P, 2\epsilon)$. \square

B.1.5. Theorem 5.6

PROOF. All LPs in the returned set are of the form $LP_{MEAE}(U, P, \frac{1}{\epsilon})$. From Lemma 5.5, all strategy profiles that are feasible for these LPs are multiple ϵ -approximate Nash equilibrium with $\epsilon = \frac{2}{e}$ and hence the first condition is satisfied. From Lemma 5.4, all feasible solutions are ϵ -approximate Nash equilibria with $\epsilon = \frac{1}{2e}$ for the given game and hence the second condition is satisfied. For computing the runtime of the algorithm, we observe that the **for** loop in the algorithm runs $(e + 1)^f$ times and each time it outputs an LP of size $2mf + 2m + 2$ which can be appended to a data structure such as a list in constant time. Thus, the algorithm can be implemented in time $O((e + 1)^f(2mf + 2m + 2))$. \square

B.2. Proofs for Section 5.2

B.2.1. Lemma 5.9. For arbitrary real numbers, we have, by matching coefficients on LHS and RHS,

$$\sum_{a \in A} \prod_{i \in [n]} x_{a_i}^i \equiv \prod_{i \in [n]} \sum_{l \in [m]} x_l^i \quad (18)$$

From (9), by combining the two product terms into one, we get: $u_j(\sigma) = \sum_{a \in A} \prod_{i \in [n]} \sigma_{a_i}^i \alpha_{a_i}^{i,j}$. From (18), with $x_{a_i}^i = \sigma_{a_i}^i \alpha_{a_i}^{i,j}$, we get

$$u_j(\sigma) = \prod_{i \in [n]} \sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j} \quad (19)$$

This function is not a function of j 's strategy. When player j 's strategy is pure and its support is just an action, say $l \in [m]$, substituting value of σ_j as e_l in (19), we have:

$$u_j(e_l, \sigma_{-j}) = \alpha_l^{j,j} \prod_{i \in [n] \setminus \{j\}} \sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j} = \alpha_l^{j,j} u'_{-j}(\sigma) \quad (20)$$

A necessary and sufficient condition for Nash equilibrium is that only the best pure responses can be in support of each player's strategy. Let $\sigma = (\sigma^1, \sigma^2, \dots, \sigma^n)$ be a Nash equilibrium for the above game. Let E be the set of all Nash equilibria for the given game. Therefore, for any player, j , with a positive payoff and any action a in support of σ^j , we have:

$$\sigma \in E \iff u_j(e_a, \sigma_{-a}) \geq u_j(e_l, \sigma_{-a}), \forall l \in [m] \quad (21)$$

$$\iff \alpha_a^{j,j} u'_{-j}(\sigma) \geq \alpha_l^{j,j} u'_{-j}(\sigma) \quad (\text{Substituting from (20)}) \quad (22)$$

Assuming that $u'_{-j} > 0$ (otherwise, j can play any action without affecting his payoff, which remains 0), we have:

$$(22) \iff \alpha_a^{j,j} \geq \alpha_l^{j,j} \iff a \in S_j \quad (\text{From defn. of } S_j) \quad (23)$$

Since the above is true $\forall a \in \text{support}(\sigma^j)$ and $\forall j \in [n]$, the claim follows.

B.2.2. Lemma 5.10

PROOF. A necessary and sufficient condition for well-supported relative ϵ -approximate Nash equilibrium is that only the approximate pure best responses can be in support of each player's strategy. Let W be the set of all well-supported relative ϵ -approximate Nash equilibria (with non-zero payoffs for each player) for the given game. Therefore, for any player, j , with a positive payoff and any strategy i in support

of σ^j , we have:

$$\begin{aligned} \sigma \in W &\iff u_j(i, \sigma_{-i}) \geq (1 - \epsilon)u_l(l, \sigma_{-i}), \forall l \in [m] \\ &\iff \alpha_i^{j,j} u'_j(\sigma) \geq (1 - \epsilon)\alpha_l^{j,j} u'_j(\sigma), \forall l \in [m] \\ &\iff \alpha_i^{j,j} \geq (1 - \epsilon)\alpha_l^{j,j}, \forall l \in [m] \\ &\iff \alpha_i^{j,j} \geq (1 - \epsilon)(\max \alpha^{j,j}) \end{aligned}$$

Thus, equation (10) above is the complete characterization of well-supported relative ϵ -approximate Nash equilibria for the given game. \square

B.2.3. Theorem 5.11

PROOF. Let σ be a strategy profile for the given game for which equation (11) holds. From the definition of well-supported multiple relative ϵ -approximate Nash equilibrium (Definition 4.4), a strategy profile is well-supported multiple relative ϵ -approximate Nash equilibrium iff it is a well-supported ϵ -approximate Nash equilibrium for each constituent game of the SGM. Since equation (11) holds for σ for all constituent games, from Lemma 5.10, σ is well-supported multiple relative ϵ -approximate Nash equilibrium for the given game. Thus, from Lemma 5.10, σ satisfies equation (10) for all constituent games. Thus, σ satisfies equation (11). \square

B.3. Proofs for Section 5.3

B.3.1. Lemma 5.16

PROOF. From definition of well-supported ϵ -approximate Nash equilibrium, for any action $a \in \text{Support}(\sigma^j)$:

$$u_j(e_a, \sigma_{-j}) \geq u_j(e_l, \sigma_{-j}) - \epsilon, \forall l \in [m] \quad (24)$$

Now, we construct a t -uniform strategy profile, σ' , from σ as follows. Let $\sigma'_i = \frac{\lfloor \sigma_i t \rfloor}{t}$. The above rounding procedure can make $\|\sigma'\|_1$ greater than 1. To counter this, select any element in σ' arbitrarily (without replacement) and round it down to $\frac{\lfloor \sigma_i t \rfloor}{t}$. Repeat this until $\|\sigma'\|_1 = 1$. The above procedure is guaranteed to give a t -uniform strategy profile σ' such that:

$$|\sigma'_i - \sigma_i| \leq \frac{1}{t} \quad (25)$$

From the definition of multiplayer games of low rank, the payoff is given by:

$$u_j(\sigma) = \sum_{k=1}^K \prod_{i \in [n]} \sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j,k} \quad (26)$$

From above, when player j plays action a with probability 1 and the rest of the players play their respective strategies in σ , we have:

$$u_j(e_a, \sigma_{-j}) = \sum_{k=1}^K \alpha_a^{j,j,k} \prod_{i \in [n] \setminus \{j\}} \sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j,k} \quad (27)$$

Thus, we have:

$$\begin{aligned} |u_j(e_a, \sigma_{-j}) - u_j(e_a, \sigma'_{-j})| &= \left| \sum_{k=1}^K \alpha_a^{j,j,k} \prod_{i \in [n] \setminus \{j\}} \sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j,k} \right. \\ &\quad \left. - \sum_{k=1}^K \alpha_a^{j,j,k} \prod_{i \in [n] \setminus \{j\}} \sum_{l \in [m]} \sigma'_l{}^i \alpha_l^{i,j,k} \right| \end{aligned}$$

By taking the outermost summation and the common factor $\alpha_a^{j,j,k}$ out, we get:

$$\begin{aligned} &|u_j(e_a, \sigma_{-j}) - u_j(e_a, \sigma'_{-j})| \\ &= \left| \sum_{k=1}^K \alpha_a^{j,j,k} \left(\prod_{i \in [n] \setminus \{j\}} \sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j,k} - \prod_{i \in [n] \setminus \{j\}} \sum_{l \in [m]} \sigma'_l{}^i \alpha_l^{i,j,k} \right) \right| \\ &= \left| \sum_{k=1}^K \alpha_a^{j,j,k} \left(\prod_{i \in [n] \setminus \{j\}} x_{i,k} - \prod_{i \in [n] \setminus \{j\}} x'_{i,k} \right) \right| \end{aligned} \quad (28)$$

Where, $x_{i,k} = \sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j,k} = (\alpha^{i,j,k})^T \sigma$ and

$x'_{i,k} = \sum_{l \in [m]} \sigma'_l{}^i \alpha_l^{i,j,k} = (\alpha^{i,j,k})^T \sigma'$. From Lemma 5.14 and Equation 25, we have $|x_{i,k} - x'_{i,k}| \leq \frac{m}{t}$. From Lemma 5.15 and the above, we have:

$$\left| \prod_{i \in [n] \setminus \{j\}} x_{i,k} - \prod_{i \in [n] \setminus \{j\}} x'_{i,k} \right| \leq \frac{(n-1)m}{t}$$

From above and Equation 28, we have:

$$\begin{aligned} |u_j(e_a, \sigma_{-j}) - u_j(e_a, \sigma'_{-j})| &\leq \left| \sum_{k=1}^K \alpha_a^{j,j,k} \frac{(n-1)m}{t} \right| \\ &\leq \frac{(n-1)mk}{t} \because \alpha_a^{j,j,k} \leq 1 \end{aligned} \quad (29)$$

From above and Equation 24, we have, for every action a in support of σ' :

$$u_j(e_a, \sigma'_{-j}) \geq u_j(e_l, \sigma'_{-j}) - \frac{2(n-1)mk}{t} - \epsilon, \forall l \in [m]$$

Thus, from the definition, σ' is a well-supported $\epsilon + \frac{2(n-1)mk}{t}$ -approximate Nash equilibrium. \square

B.3.2. Theorem 5.17

PROOF. From Lemma 5.16, given a well-supported τ -approximate Nash equilibrium strategy profile, we can always construct a t -uniform strategy profile that is a well-supported $\tau + \frac{(n-1)mk}{t}$ -approximate Nash equilibrium. From the definition of well-supported multiple ϵ -approximate Nash equilibrium, σ is a well-supported τ -approximate Nash equilibrium strategy profile for each of the constituent games. The construction of σ' from σ , as described in the proof of Lemma 5.16, is independent of payoff functions for constituent games and hence the lemma applies simultaneously to all the constituent games of the SGM. Hence σ' is a well-supported $\tau + \frac{(n-1)mk}{t}$ -approximate Nash equilibrium for all the constituent games of the SGM. Thus, from

the definition of well-supported multiple ϵ -approximate Nash equilibrium, σ' is a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau + \frac{(n-1)mk}{t}$ and the claim follows. \square

B.4. Proofs for Section 6.2

B.4.1. Lemma 6.2

PROOF. Let $\sigma = (\sigma^1, \dots, \sigma^n)$ be a Nash equilibrium. Construct a multiset M^j by sampling independently at random from the set of actions according to distribution σ^j . Construct strategy ρ^j by assigning probability of $\frac{l}{t}$ to an action that appears l times in M_j and $\rho^{-j} = \sigma^{-j}$. We prove that ρ is an $\tau + \delta$ -well supported multiple payoff approximate Nash equilibrium **w.p.p.**

Let $M_s^j \in [m]$ be a random variable that denotes the s^{th} element of M^j . Let $Y_s = u_l^i(e_{M_s^j}, e_k, \sigma^{-j,l})$ be a random variable that is equal to payoff for player l , when player l plays action k and player j plays action M_s^j in constituent game i . Let μ be the mean of random variables $\{Y_1, \dots, Y_t\}$. Then, we have:

$$\begin{aligned} \mu &= \sum_{s=1}^t \frac{u_l^i(e_{M_s^j}, e_k, \sigma^{-j,l})}{t} \\ &= \sum_{s=1}^t \frac{u_l^i(e_{M_s^j}, e_k, \rho^{-j,l})}{t} \quad (\mathbf{from} \ (\rho^{-j} = \sigma^{-j})) \\ &= u_l^i\left(\sum_{s=1}^t \frac{e_{M_s^j}}{t}, e_k, \rho^{-j,l}\right) \quad (\mathbf{from} \ (\text{definition of } u_l^i)) \\ &= u_l^i(e_k, \rho^{-l}) \end{aligned}$$

Where, the last equality follows from the fact that if an action occurs l times in M^j , it gets a weight of $\frac{l}{t}$ is ρ^j . Also, from construction of ρ , it follows that $E(\rho_k^j) = \sigma_k^j$. From linearity of expectation it follows that $E(u_l^i(e_k, \rho^{-l})) = u_l^i(e_k, \sigma^{-l})$. Thus, $\forall i \in [f], l \in [n] - \{j\}, k \in [m]$, $u_l^i(e_k, \rho^{-l})$ is the mean of t i.i.d. random variables, each with expectation $u_l^i(e_k, \sigma^{-l})$. Payoffs are in $[0, 1]$ and the same bound applies to all these random variables. Also, since only player j 's strategy is changed, $u_j^i(e_k, \sigma^{-j}) = u_j^i(e_k, \rho^{-j})$.

Let $A(i, k, l)$ be the event $|u_l^i(e_k, \rho^{-l}) - u_l^i(e_k, \sigma^{-l})| \geq \frac{\delta}{2}$. From Hoeffding inequality, we have:

$$Pr[A(i, k, l)] \leq 2 \exp\left(\frac{-t\delta^2}{2}\right)$$

Let $A = \bigcup_{i \in [f], k \in [m], l \in [n] - \{j\}} A(i, k, l)$. From union bound:

$$Pr[A] \leq 2fm(n-1) \exp\left(\frac{-t\delta^2}{2}\right)$$

Thus, $2fmn \exp\left(\frac{-t\delta^2}{2}\right) = 1 \implies Pr[A^c] > 0$. Therefore, $t = \frac{2 \log(2fmn)}{\delta^2}$ ensures that $\exists \rho$ for which event A^c occurs. Let π be the strategy profile for which this happens. From definition of A^c and well-supported multiple ϵ -approximate Nash equilibrium, we have, $\forall i \in [f], \forall k \in [m], \forall l \in [n]$:

$$\begin{aligned} u_l^i(e_k, \sigma^{-l}) &\leq u_l^i(e_k, \pi^{-l}) + \tau \quad (\mathbf{from} \ (\text{Definition 4.3})) \\ \implies u_l^i(e_k, \pi^{-l}) &\leq u_l^i(e_k, \sigma^{-l}) + \tau + \delta \quad (\mathbf{from} \ (\text{definition of } A^c)) \end{aligned}$$

Thus π is well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau + \delta$. \square

B.4.2. Theorem 6.3

PROOF. Let $t = \frac{2n^2 \log(2fm(n-1))}{\delta^2}$. Then by application of Lemma 6.2 for player 1, we get a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau + \frac{\delta}{n}$ with t -uniform strategy of player 1. To the resulting strategy profile, we can apply Lemma 6.2 to the strategy for player 2 and get a strategy profile that is a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau + \frac{2\delta}{n}$ with t -uniform strategy of players 1 and 2. We can do this successively for all players and get a t -uniform strategy profile which is a well-supported multiple ϵ -approximate Nash equilibrium with $\epsilon = \tau + \delta$ \square