Function Approximation for Solving Stackelberg Equilibrium in Large Perfect Information Games

Chun Kai Ling, J. Zico Kolter, Fei Fang

School of Computer Science, Carnegie Mellon University chunkail@cs.cmu.edu, zkolter@cs.cmu.edu, feif@cs.cmu.edu

Abstract

Function approximation (FA) has been a critical component in solving large zero-sum games. Yet, little attention has been given towards FA in solving general-sum extensiveform games, despite them being widely regarded as being computationally more challenging than their fully competitive or cooperative counterparts. A key challenge is that for many equilibria in general-sum games, no simple analogue to the state value function used in Markov Decision Processes and zero-sum games exists. In this paper, we propose learning the Enforceable Payoff Frontier (EPF)-a generalization of the state value function for general-sum games. We approximate the optimal Stackelberg extensive-form correlated equilibrium by representing EPFs with neural networks and training them by using appropriate backup operations and loss functions. This is the first method that applies FA to the Stackelberg setting, allowing us to scale to much larger games while still enjoying performance guarantees based on FA error. Additionally, our proposed method guarantees incentive compatibility and is easy to evaluate without having to depend on self-play or approximate best-response oracles.

1 Introduction

A central challenge in modern game solving is to handle large game trees, particularly those too large to traverse or even specify. These include board games like Chess, Poker (Silver et al. 2018, 2016; Brown and Sandholm 2017, 2019; Moravčík et al. 2017; Bakhtin et al. 2021; Gray et al. 2021) and modern video games with large state and action spaces (Vinyals et al. 2019). Today, scalable game solving is frequently achieved via function approximation (FA), typically by using neural networks to model state values and harnessing the network's ability to generalize its evaluation to states never encountered before (Silver et al. 2016, 2018; Moravčík et al. 2017; Schmid et al. 2021). Methods employing FA have achieved not only state-of-the-art performance, but also exhibit more human-like behavior (Kasparov 2018).

Surprisingly, FA is rarely applied to solution concepts used in general-sum games such as Stackelberg equilibrium, which are generally regarded as being more difficult to solve than the perfectly cooperative/competitive Nash equilibrium. Indeed, the bulk of existing literature centers around on methods such as exact backward induction (Bosanský et al. 2015; Bošanskỳ et al. 2017), incremental strategy generation (Černỳ, Bošanskỳ, and Kiekintveld 2018; Cermak et al. 2016; Karwowski and Mańdziuk 2020), and mathematical programming (Bosansky and Cermak 2015).¹ While exact, these methods rarely scale to large game trees, especially those too large to traverse, severely limiting our ability to tackle general-sum games that are of practical interest, such as those in security domains like wildlife poaching prevention (Fang et al. 2017) and airport patrols (Pita et al. 2008). For non-Nash equilibrium in general-sum games, the value of a state often *cannot* be summarized as a scalar (or fixed sized vector), rendering the direct application of FA-based zero-sum solvers like (Silver et al. 2018) infeasible.

In this paper, we propose applying FA to model the Enforceable Payoff Frontier (EPF) for each state and using it to solve for the Stackelberg extensive-form correlated equilibrium (SEFCE) in two-player games of perfect information. Introduced in (Bošanský et al. 2017; Bosanský et al. 2015; Letchford and Conitzer 2010), EPFs capture the tradeoff between player payoffs and is analogous to the state value in zero-sum games.² Specifically, we (i) study the pitfalls that can occur with using FA in general-sum games, (ii) propose a method for solving SEFCEs by modeling EPFs using neural networks and minimizing an appropriately designed Bellman-like loss, and (iii) provide guarantees on incentive compatibility and performance of our method. Our approach is the first application of FA in Stackelberg settings without relying on best-response oracles for performance guarantees. Experimental results show that our method can (a) approximate solutions in games too large to explicitly traverse, and (b) generalize learned EPFs over states in a repeated setting where game payoffs vary based on features.

2 Preliminaries and Notation

A 2-player *perfect information* game G is represented by a finite game tree with game states $s \in S$ given by vertices and action space $\mathcal{A}(s)$ given by directed edges starting from s.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Meta-game solving (Lanctot et al. 2017; Wang et al. 2019) is used in zero-sum games, but not general-sum Stackelberg games.

²The idea of an EPF was initially used by (Letchford and Conitzer 2010) to give a polynomial time solution for SSEs. However, they (as well as other work) do not propose any naming.

Each state belongs to either player P_1 or P_2 ; we denote these disjoint sets by S_1 and S_2 respectively. Every leaf (terminal state) $\ell \in \mathcal{L} \subseteq S$ of G is associated with payoffs, given by $r_i(\ell)$ for each player *i*. Taking action $a \in \mathcal{A}(s)$ at state $s \notin \mathcal{L}$ leads to s' = T(a; s), where $s' \in S$ is the next state and Tis the deterministic transition function. Let $C(s) = \{s' \mid T(a; s) = s', a \in \mathcal{A}(s)\}$ denote the immediate children of *s*. We say that state *s* precedes (\Box) state *s'* if $s \neq s'$ and *s* is an ancestor of *s'* in G, and write \Box if allowing s = s'. An action $a \in \mathcal{A}(s)$ leads to *s'* if $s \sqsubset s'$ and $T(a; s) \sqsubseteq s'$. With a slight abuse of notation, we denote $T(a; s) \sqsubseteq s'$ by $a \sqsubset s'$ or $(s, a) \sqsubset s'$. Since G is a tree, for states *s*, *s'* where $s \sqsubset s'$, exactly one $a \in \mathcal{A}(s)$ such that $(s, a) \sqsubseteq s'$. We use the notation $\sqsupseteq and \sqsupset$ when the relationships are reversed.

A strategy π_i , where $i \in \{P_1, P_2\}$, is a mapping from state $s \in S_i$ to a distribution over actions $\mathcal{A}(s)$, i.e., $\sum_{a \in \mathcal{A}(s)} \pi_i(a; s) = 1$. Given strategies π_1 and π_2 , the probability of reaching $\ell \in \mathcal{L}$ starting from s is given by $p(\ell|s; \pi_1, \pi_2) = \prod_{i \in \{P_1, P_2\}} \prod_{(s', a); s \sqsubseteq s', (s', a) \sqsubseteq \ell, s' \in S_i} \pi_i(a; s')$, and player i's expected payoff starting from s is $R_i(s; \pi_1, \pi_2) = \sum_{\ell \in \mathcal{L}} p(\ell|s; \pi_1, \pi_2) r_i(\ell)$. We use as shorthand $p(\ell; \pi_1, \pi_2)$ and $R_i(\pi_1, \pi_2)$ if s is the root. A strategy π_2 is a *best response* to a strategy π_1 if $R_2(\pi_1, \pi_2) \ge R_2(\pi_1, \pi'_2)$ for all strategies π'_2 . The set of best responses to π_1 is written as $BRS_2(\pi_1)$.

The grim strategy $\arg \min_{\pi_1} \max_{\pi_2} R_2(\pi_1, \pi_2)$ of P₁ towards P₂ is one which guarantees the lowest payoff for P₂. Conversely, the *joint altruistic strategy* $\arg \max_{\pi_1,\pi_2} R_2(\pi_1,\pi_2)$ is one which maximizes P₂'s payoff. We restrict grim and altruistic strategies to those which are subgame-perfect, i.e., they remain the optimal if the game was rooted at some other state.³ Grim and altruistic strategies ignore P₁'s own payoffs and can be computed by backward induction. For each state, we denote by $\underline{V}(s)$ and $\overline{V}(s)$ the internal values of P₂ for grim and altruistic strategies obtained via backward induction.

2.1 Stackelberg Equilibrium in Perfect Information Games

In a Strong Stackelberg equilibrium (SSE), there is a distinguished *leader* and *follower*, which we assume are P₁ and P₂ respectively. The leader *commits* to any strategy π_1 and the follower best responds to the leader, breaking ties by selecting $\pi_2 \in BRS_2(\pi_1)$ such as to benefit the leader. ⁴ Solving for the SSE entails finding the optimal commitment for the *leader*, i.e., a pair $\pi = (\pi_1, \pi_2)$ such that $\pi_2 \in BRS_2(\pi_1)$ and $R_1(\pi_1, \pi_2)$ is to be maximized.

It is well-known that the optimal SSE will perform no worse (for the leader) than Nash equilibrium, and often much better. Consider the game in Figure 1a with $k_1 = k_2 =$ 0. If the expected follower payoff from staying is less than 0, then it would exit immediately. Hence, solutions such as the subgame perfect Nash gives a leader payoff of 0. The optimal Stackelberg solution is for the leader to commit to a uniform strategy—this ensures that staying yields the follower a payoff of 0, which under the tie-breaking assumptions of SSE nets the leader a payoff of 4.5.

Stackelberg Extensive-Form Correlated Equilibrium For this paper, we will focus on a relaxation of the SSE known as the Stackelberg extensive-form correlated equilibirum (SEFCE), which allows the leader to explicitly recommend actions to the follower *at the time of decision making*. If the follower deviates from the recommendation, the leader is free to retaliate—typically with the grim strategy. In a SEFCE, P₁ takes and recommends actions to maximize its reward, subject to the constraints that the recommendations are sufficiently appealing to P₂ relative to threat of P₂ facing the grim strategy after any potential deviation.

Definition 1 (Minimum required incentives). Given $s \in S_2$, $s' \in C(s)$, we define the minimum required incentive $\tau(s') = \max_{s' \in C(s); s' \neq s'} \underline{V}(s')$, i.e., the minimum amount that P_1 needs to promise P_2 under s' for it to be reached.

Definition 2 (SEFCE). A strategy pair $\pi = (\pi_1, \pi_2)$ is a SEFCE if it is incentive compatible, i.e., for all $s \in S_2$, $a \in \mathcal{A}(s)$, $\pi_2(a; s) > 0 \Longrightarrow R_2(T(a; s); \pi_1, \pi_2) \ge \tau(T(a; s))$. Additionally, π is optimal if $R_1(\pi_1, \pi_2)$ is maximized.

In Section 3, we describe how optimal SEFCE can be computed in polynomial time for perfect information games.

2.2 Function Approximation of State Values

When finding Nash equilibrium in perfect information games, the value v_s of a state is a crucial quantity which summarizes the utility obtained from s onward, assuming optimal play from all players. It contains sufficient information for one to obtain an optimal solution after using them to 'replace' subtrees. Typically v_s should only rely on states $s' \supseteq s$. In zero-sum games, $v_s = V_s$ while in cooperative games, $v_s = \overline{V}_s$. Knowing the true value of each state immediately enables the optimal policy via one-step lookahead. While v_s can be computed over all states by backward induction, this is not feasible when G is large. A standard workaround is to replace v_s with an approximate \tilde{v}_s which is then used in tandem with some search algorithm (depthlimited search, Monte-Carlo tree search, etc.) to obtain an approximate solution. Today, \tilde{v}_s is often *learned*. By representing \tilde{v} with a rich function class over state features (typically using a neural network), modern solvers are able to generalize \tilde{v} across large state spaces without explicitly visiting every state, thus scaling to much larger games.

Fitted Value Iteration. A class of methods closely related to ours is Fitted Value Iteration (FVI) (Lagoudakis and Parr 2003; Dietterich and Wang 2001; Munos and Szepesvári 2008). The idea behind FVI is to optimize for parameters such as to minimize the *Bellman loss* over sampled states by treating it as a regular regression problem. ⁵ Here, the

³This is to avoid strategies which play arbitrarily at states which have 0 probability of being reached.

⁴Commitment rights are justified by repeated interactions. If the P_1 reneges on its commitment, P_2 plays another best response, which is detrimental to the leader. This setting is unlike (De Jonge and Zhang 2020) which uses binding agreements.

⁵We distinguish RL and FVI in that the transition function is known explicitly and made used of in FVI.



Figure 1: (a) Game tree to illustrate computation of SEFCE. Leader \bigcirc , follower \bigcirc and leaf \square states are vertices and edges are actions. (b-d) EPFs at s', after exiting and s. The x and y axes are follower (μ_2) and leader payoffs ($U_s(\mu)$). In (b) the pink regions give P₂ too little and are truncated. In (d), the pink regions are not part of the upper concave envelope and removed.

Bellman loss measures the distance between \tilde{v}_s and the estimated value using one-step lookahead using \tilde{v} . If this distance is 0 for all s, then \tilde{v} matches the optimal v. In practice, small errors in FA accumulate and cascade across states, lowering performance. Thus, it is important to bound performance as a function of the Bellman loss over all s.

2.3 Related Work

Some work has been done in generalizing state values in general-sum games, but few involve learning them. Related to ours is (Murray and Gordon 2007; MacDermed et al. 2011; Dermed and Charles 2013), which approximate the achievable set of payoffs for correlated equilibrium, and eventually SSE (Letchford et al. 2012) in stochastic games. These methods are analytical in nature and scale poorly. (Pérolat et al. 2017; Greenwald et al. 2003) propose a Qlearning-like algorithm over general-sum Markov games, but do not apply FA and only consider stationary strategies which preclude strategies involving long range threats like the SSE. (Zinkevich, Greenwald, and Littman 2005) show a class of general-sum Markov games where value-iteration like methods will necessarily fail. (Zhong et al. 2021) study reinforcement learning in the Stackelberg setting, but only consider followers with myopic best responses. (Castelletti, Pianosi, and Restelli 2011) apply FVI in a multiobjective setting, but do not consider the issue of incentive compatibility. Another approach is to apply reinforcement learning and self-play (Leibo et al. 2017). Recent methods account for the nonstationary environment each player faces during training (Foerster et al. 2017; Perolat et al. 2022); however they have little game theoretical guarantees in terms of incentive compatibility, particularly in non zero-sum games.

3 Review: Solving SEFCE via Enforceable Payoff Frontiers

In Section 2, we emphasized the importance of the value function v in solving zero-sum games. In this section, we review the analogue for SEFCE in the general-sum games, which we term as *Enforceable Payoff Frontiers* (EPF). Introduced in (Letchford and Conitzer 2010), the EPF at state s is a *function* $U_s : \mathbb{R} \to \mathbb{R} \cup \{-\infty\}$, such that $U_s(\mu_2)$ gives the maximum leader payoff for a SEFCE for a game rooted at s, on condition that P_2 obtains a payoff of μ_2 . All leaves $s \in \mathcal{L}$

have degenerate EPFs $U_s(r_2(s)) = r_1(s)$ and $-\infty$ everywhere else. EPFs capture the tradeoff in payoffs between P₁ and P₂, making them useful for solving SEFCEs. We now review the two-phase algorithm of (Bošanskỳ et al. 2017) using the example game in Figure 1a with $k_1 = k_2 = 0$. This approach forms the basis for our proposed FA method.

Phase 1: Computing EPF by Backward Induction. The EPF at s' is given by the line segment connecting payoffs of its children EPF and $-\infty$ everywhere else. This is because the leader is able to freely mix over actions. To compute U_s , we consider in turn the EPFs after staying or exiting. Case 1: P_1 is recommending P_2 to stay. For incentive compatibility, it needs to promise P_2 a payoff of at least 0 under T(stay; s) = s'. Thus, we **left-truncate** the regions of the EPF at s' which violate this promise, leaving behind the blue segment (Figure 1b), which represents the payoffs at s' that are *enforceable* by P_1 . Case 2: P_1 is recommending P_2 to exit. To discourage P_2 from staying, it commits to the grim strategy at s' if P₂ chooses to stay instead, yielding P₂ a payoff of $-1 \le k_2 = 0$. Hence, no truncation is needed and the set of enforceable payoffs is the (degenerate) blue line segment (Figure 1c). Finally, to recover U_s , observe that we can achieve payoffs on any line segment connecting point across the EPFs of s's children. This union of points on such lines (ignoring those leader-dominated) is given by the upper concave envelope of the blue segments in Figure 1b and 1c; this removes $\{(0,0)\}$, giving the EPF in Figure 1d.

More generally, let g_1 and g_2 be functions such that $g_j : \mathbb{R} \to \mathbb{R} \cup \{-\infty\}$. We denote by $g_1 \wedge g_2$ their upper concave envelope, i.e., $\inf\{h(\mu) \mid h \text{ is concave and } h \ge \max\{g_1, g_2\} \text{ over } \mathbb{R}\}$. Since \wedge is associative and commutative, we use as shorthand $\wedge_{\{\cdot\}}$ when applying \wedge repeatedly over a finite set of functions. In addition, we denote $g \triangleright t$ as the left-truncation of the g with threshold $t \in \mathbb{R}$, i.e., $[g \triangleright t](\mu) = g(\mu)$ if $\mu \ge t$ and $-\infty$ otherwise. Note that both \wedge and \triangleright are closed over concave functions. For any $s \in S$, its EPF U_s can be concisely written in terms of its children EPF $U_{s'}$ (where $s' \in C(s)$) using \wedge , \triangleright and $\tau(s')$.

$$U_{s}(\mu) = \begin{cases} \left[\bigwedge_{s' \in \mathcal{C}(s)} U_{s'} \right](\mu) & \text{if } s \in \mathcal{S}_{1} \\ \left[\bigwedge_{s' \in \mathcal{C}(s)} U_{s'} \triangleright \tau(s') \right](\mu) & \text{if } s \in \mathcal{S}_{2} \end{cases}, \quad (1)$$

which we apply in a bottom-up fashion to complete Phase 1.

Phase 2: Extracting Strategies from EPF. Once U_s has been computed for all $s \in S$, we can recover the optimal strategy π_1 by applying one-step lookahead starting from the root. First, we extract (OPT_2, OPT_1) , the coordinates of the maximum point in U_{root} , which contain payoffs under the optimal π . Here, this is (0, 4.5). We initialize $\mu_2 = \mathsf{OPT}_2$, which represents P_1 's promised payoff to P_2 at the current state s. Next, we traverse G depth-first. By construction, $U_s(\mu_2) > -\infty$ and the point $(\mu_2, U_s(\mu_2))$ is the convex combination of either 1 or 2 points belonging to its children EPFs. The mixing factors correspond to the optimal strategy $\pi(a; s)$. If there are 2 distinct children s', s'' with mixing factor α', α'' , we repeat this process for s', s'' with $\mu_2' = \mu_2/\alpha', \mu_2'' = \mu_2/\alpha''$, otherwise we repeat the process for s' and $\mu'_2 = \mu_2$. For our example, we start at $s, \mu_2 = 0$, which was obtained by P2 playing 'stay' exclusively, so we keep μ_2 and move to s'. At s', $\mu = 0$ by mixing uniformly, which gives us the result in Section 2.

Theorem 1 ((Bošanský et al. 2017; Bosanský et al. 2015)). (i) U_s is piecewise linear concave with number of knots⁶ no greater than the number of leaves beneath s. (ii) Using backward induction, SEFCEs can be computed in polynomial time (in |S|) even in games with chance. EPFs continue to be piecewise linear concave.

Markovian Property. Just like state values v_s in zero-sum games, we can replace any internal vertex s in G with its EPF while not affecting the optimal strategy in all other branches of the game. This can done by adding a single leader vertex with actions leading to terminal states with payoffs corresponding to the knots of U_s . Since U_s is obtained via backward induction, it only depends on states beneath s. In fact, if two games G and G' (which could be equal to G) shared a common subgame rooted in s and s' respectively, we could reuse the U_s found in G for $U_{s'}$ in G'. This observation underpins the inspiration for our work—if s and s' are similar in some features, then U_s and $U_{s'}$ are likely similar and it should be possible to *learn* and *generalize* EPFs over states.

4 Challenges in Applying FA to EPF

We now return to our original problem of applying FA to find SEFCE. Our idea, outlined in Algorithm 1 and 2 is a straightforward extension of FVI. Suppose each state has features f(s)—in the simplest case this could be a state's history. We design a neural network $E_{\phi}(f)$ parameterized by ϕ . This network maps state features f(s) to some representation of \tilde{U}_s , the approximated EPFs. To achieve a good approximation, we optimize ϕ by minimizing an appropriate Bellman-like loss (over EPFs) based on Equation (1) while using our approximation \tilde{U}_s in lieu of U_s . Despite its simplicity, there remain several design considerations.

EPFs are the 'right' object to learn. Unlike state values, representing an exact EPF at a state *s* could require more than constant memory since the number of knots could be linear in the number of leaves underneath it (Theorem 1). Can we get away with summarizing a state with a scalar or a

Algorithm 1: Training Pipeline

- 1: Sample trajectory $s_{\text{new}}^{(1)}, \ldots, s_{\text{new}}^{(t)}$
- 2: Update replay buffer \mathcal{B} with $s^{(1)}, \ldots, s^{(t)}$
- 3: for $i \in \{1, ..., t\}$ do
- 4: Sample batch $S = \{s^{(1)}, \dots s^{(n)}\} \subseteq \mathcal{B}$
- 5: $\ell \leftarrow \text{COMPUTELOSS}(S; E_{\phi})$
- 6: Update ϕ using $\partial \ell / \partial \phi$

Algorithm 2: COMPUTELOSS $(S; E_{\phi})$

 $\begin{array}{ll} & \text{1: } \textbf{for } i \in \{1 \dots n\} \, \textbf{do} \\ & \text{2: } \quad \tilde{U}_{s^{(i)}} \leftarrow E_{\phi}(f(s^{(i)})) \\ & \text{3: } \quad \tilde{U}_{s^{(j)}_{\text{next}}} \leftarrow E_{\phi}(f(s^{(j)}_{\text{next}})) \quad \text{for all } s^{(j)}_{\text{next}} \in \mathcal{C}(s^{(i)}) \\ & \text{4: } \quad \text{Compute } \tilde{U}^{\text{target}}_{s^{(i)}} \text{ using Equation (1) and } \{\tilde{U}_{s^{(j)}_{\text{next}}}\} \\ & \text{5: } \textbf{return } \sum_{i} L(\tilde{U}_{s^{(i)}}, \tilde{U}^{\text{target}}_{s^{(i)}}) \end{array}$

small vector? Unfortunately, any 'lossless summary' which enjoys the Markovian property necessarily encapsulates the EPF. To see why, consider the class of games G_k in Figure 1a with $k_1 = -2$ and $k = k_2 \in [-1, 1]$. The optimal leader payoff for any G_k is $\frac{9-11k}{2}$, which is precisely $U_{s'}(k)$ (Figure 1b). Now consider any lossless summary for s' and use it to solve *every* G_k . The resultant optimal leader payoffs can recover $U_{s'}(\mu_2)$ between $\mu_2 \in [-1, 1]$. This implies that no lossless summary more compact than the EPF exists.

Unfulfillable Promises Arising from FA Error. Consider the game in Figure 2a with $k_1 = -10, k_2 = -1$. The exact $U_{s'}$ is the line segment combining the points (-1, 10) and $(1-\epsilon, -1)$, shown in green in Figure 3a. However, let us suppose that due to function approximation we instead learned the blue line segment containing (-1, 10) and (1, -1). Performing Phase 2 using \tilde{U} , the policy extracted at s' is once again the uniform policy and requires us to promise the follower a utility of 1 in s''. However, achieving a payoff of 1 is *impossible* regardless of how much the leader is willing to sacrifice, since the maximum outcome under s'' is $1 - \epsilon$. Since this is an unfulfillable promise, the follower's best responds by exiting in s, which gives the leader a payoff of -10. In general, unfulfillable promises due to small FA error can lead to arbitrarily low payoffs. In fact, one could argue that \tilde{U} does not even *define* a valid policy.

Costly Promises. Consider the case where $k_1 = -30, k_2 = 1$ while keeping $\tilde{U}_{s'}$ the same. Here, the promise of 1 at s'' is fulfillable, but involves incurring a cost of -30, which is even lower than having follower staying (Figure 3b). In general, this problem of *costly promises* stems from the EPF being wrongly estimated, even for a small range of μ_2 . We can see how costly promises arise even from small ϵ is. The underlying issue is that in general, U_s can have large Lipschitz constants (e.g., proportionate to $(\max_s r_1(s) - \min_s r_1(s)) / (\min |r_2(s) - r_2(s)|)$). The existence of costly payoffs rules out EPF representations based

⁶Knots are where the slope of the EPF changes.



(a) Game tree to illustrate unful- (b) Stage game used in the fillable and costly promises. TANTRUM game.

Figure 2: Games used in Sections 4 and 6. Leader \bigcirc , follower \bigcirc and leaf \square states are vertices and edges are actions.



Figure 3: EPFs for (a) unfulfillable promises and (b) costly promises. Blue lines are estimated EPFs $\tilde{U}_{s'}$, solid and dotted green lines are true EPFs $U_{s'}$, $U_{s''}$. In both cases, FA error leads us to believe that the payoff given by the blue square at (0, 4.5) can be achieved by mixing the endpoints of $\tilde{U}_{s'}$ with probability $\alpha' = \alpha'' = 0.5$ (black curves).

on discretizing the space of μ_2 , since small errors incurred by discretization could lead to huge drops in performance.

5 FA of EPF with Performance Guarantees

We now design our method using the insights from Section 4. We learn EPFs without relying on discretization over P_2 payoffs μ_2 . Unfulfillable promises are avoided entirely by ensuring that the set of μ_2 where $\tilde{U}_s(\mu_2) > -\infty$ lies within some known set of achievable P_2 payoffs, while costly promises are mitigated by suitable loss functions.

Representing EPFs using Neural Networks. Our proposed network architecture represents EPFs by a small set of $m \ge 2$ points $P_{\phi}(s) = \{(x_j, y_j)\}$, for $j \in [m]$. Here, m is a hyperparameter trading off complexity of the neural network E_{ϕ} with its representation power. The approximated EPF \tilde{U}_s is the linear interpolation of these m points; and $\tilde{U}_s = -\infty$ if $\mu_2 > \max_j x_j$ or $\mu_2 < \min_j x_j$. For now, we make the assumption that follower payoffs under the altruistic and grim strategy ($\overline{V}(s)$ and $\underline{V}(s)$) are known *exactly* for all states. Through the architecture of E_{ϕ} that for all $j \in [m]$, we have $\underline{V}(s) \le x_j \le \overline{V}(s)$. As we will see, this helps avoid unfulfillable promises and allows for convenient loss functions.

Concretely, our network $E_{\phi}(f(s); \underline{V}(s), V(s))$ takes in as inputs state features f(s), lower and upper bounds $\underline{V}(s) \leq \overline{V}(s)$ and outputs a matrix in $\mathbb{R}^{m \times 2}$ representing $\{(x_j, y_j)\}$ where $x_1 = \underline{V}(s)$ and $x_m = \overline{V}(s)$. For simplicity, we use a multilayer feedforward network with depth d, width w and ReLU activations for each layer. Serious applications should utilize domain specific architectures. Denoting the output of the last fully connected layer by $h^{(d)}(f(s)) \in \mathbb{R}^w$, for $j \in \{2 \dots m-1\}$ and $k \in [m]$ we set

$$\begin{aligned} x_j &= \sigma \left(z_{x,j}^T h^{(d)}(f(s)) + b_{x,j} \right) \cdot \left(\overline{V}(s) - \underline{V}(s) \right) + \underline{V}(s), \\ y_k &= z_{y,k}^T h^{(d)}(f(s)) + b_{y,k}, \end{aligned}$$

and $x_1 = \underline{V}(s)$ and $x_m = \overline{V}(s)$, where $\sigma(x) = 1/(1 + \exp(-x))$. Here, $z_{x,j}, z_{y,k} \in \mathbb{R}^w$ and $b_{x,j}, b_{y,k} \in \mathbb{R}$ are weights and biases, which alongside the parameters from feedforward network form the network parameters ϕ to be optimized. Since \tilde{U}_s is represented by its knots (given by $P_{\phi}(s)$), Λ and consequently, (1) may be performed *explicitly* and *efficiently*, returning an entire EPF represented by its knots (as opposed to the EPF evaluated at a single point). This is crucial, since the computation is performed every state every iteration (Line 4 of Algorithm 2).

Loss Functions for Learning EPFs. Given 2 EPFs $\tilde{U}_s, \tilde{U}'_s$ we minimize the following loss to mitigate costly promises,

$$L_{\infty}(\tilde{U}_{s}, \tilde{U}_{s}') = \max_{\mu_{2}} |\tilde{U}_{s}(\mu_{2}) - \tilde{U}_{s}'(\mu_{2})|.$$

 L_{∞} was chosen specifically to incur a large loss if the approximation is wildly inaccurate in a small range of μ_2 (e.g., Figure 3b). Achieving a small loss requires that $\tilde{U}_s(\mu_2)$ approximates $\tilde{U}'_s(\mu_2)$) well for all μ_2 . This design decision is particularly important. For example, contrast L_{∞} with another intuitive loss $L_2(\tilde{U}_s, \tilde{U}'_s) = \int_{\mu_2} (\tilde{U}_s(\mu_2) - \tilde{U}'_s(\mu_2))^2 d\mu_2$. Observe that L_2 is exceedingly small in the example of Figure 3b — in fact, when ϵ is small enough leads to almost no loss, even though the policy as discussed in Section 4 is highly suboptimal. This phenomena leads to costly promises, which was indeed observed in our tests.

Our Guarantees. Any learned \tilde{U} implicitly defines a policy $\tilde{\pi}$ by one-step lookahead using Equation (1) and the method described in Phase 2 (Section 3). Extracting $\tilde{\pi}$ need not be done offline for all $s \in S$; in fact, when G is too large it is necessary that we only extract $\tilde{\pi}(\cdot; s)$ on-demand. Nonetheless, $\tilde{\pi}$ enjoys some important properties.

Theorem 2 (Incentive Compatibility). For any policy $\tilde{\pi}$ obtained using our method, any $s \in S_2$ and $a \in \mathcal{A}(s)$, we have $\tilde{\pi}_s(a; s) > 0 \Longrightarrow R_2(T(a; s); \tilde{\pi}) \ge \tau(T(a; s)).$

Theorem 3 (FA Error). If $L_{\infty}(\tilde{U}_s, \tilde{U}_s^{target}) \leq \epsilon$ for all $s \in S$, then $|R_1(\tilde{\pi}) - R_1(\pi^*)| = \mathcal{O}(D\epsilon)$ where D is the depth of G and π^* is the optimal strategy.

Here, T(a; s) is transition function (Section 2). Recall from Section 2 that for π to be an optimal SEFCE, we require (i) incentive compatibility and (ii) $R_1(\pi)$ to be maximized. Theorems 2 and 3 illustrate how our approach disentangles these criteria. Theorem 2 guarantees that P₂ will always be incentivized to follow P₁'s recommendations, i.e., there will be no unexpected outcomes arising from unfulfillable promises. Crucially, this is a hard constraint which is satisfied solely due to our choice of network architecture, which ensures that $\tilde{U}_s(\mu_2) = -\infty$ when $\mu_2 > \overline{V}_s$ for any $\tilde{\pi}$ obtained from \tilde{U} . Conversely, Theorem 3 shows that the goal of maximizing R_1 subject to incentive compatibility is achieved by attaining a small FA error across all states. This distinction is important. Most notably, incentive compatibility is no longer dependent on convergence during training. This *explicit* guarantee stands in contrast with methods employing self-play reinforcement learning agents; there, incentive compatibility follows *implicitly* from the apparent convergence of a player's strategy. This guarantee has practical implications, for example, evaluating the quality of $\tilde{\pi}$ can be done by estimating $R_1(\tilde{\pi})$ based on sampled trajectories, while implicit guarantees requires incentive compatibility to be demonstrated using some approximate best-response or acle and usually involves expensive training of a RL agent.

The primary limitation of our method is when \overline{V} and \underline{V} (and hence τ) are not known exactly. As it turns out, we can instead use upper and lower approximations while still retaining incentive compatibility. Let $\tilde{\pi}_1^{\text{grim}}$ be an *approximate grim strategy*. Define V(s) to be the expected follower payoffs at *s* when faced best-responding to $\tilde{\pi}_1^{\text{grim}}$, i.e., $R_2(s; \tilde{\pi}_1^{\text{grim}}, \pi_2)$, where $\pi_2 \in BRS_2(\tilde{\pi}_1^{\text{grim}})$. Following Definition 1, the *approximate minimum required incentive* is $\tilde{\tau}(s') = \max_{s^! \in \mathcal{C}(s); s^! \neq s'} V(s^!)$ for all $s \in S_2$, $s' \in \mathcal{C}(s)$. Similarly, let $\tilde{\pi}^{\text{alt}}$ be an *approximate joint altruistic strategy* and its resultant internal payoffs in each state be $\tilde{V}(s)$.

Under the mild assumption that $\tilde{\pi}^{\text{alt}}$ always benefits P_2 more than the $\tilde{\pi}_1^{\text{grim}}$, i.e., $\tilde{V}(s) \geq \tilde{V}(s)$ for all s, we can replace the τ , \underline{V} and \overline{V} with $\tilde{\tau}$, V and \tilde{V} and maintain incentive compatibility (Theorem 2). The intuition is straightforward: if P_2 's threats are 'good enough', parts of the EPF will still be enforceable. Furthermore, promises will always be fulfillable since EPFs domains are now limited to be no greater than $\tilde{V}(s)$, which we know can be achieved by definition. Unfortunately, Theorem 3 no longer holds, not even in terms of $\max_s |\tilde{\tau}(s) - \tau(s)|$. This is again due to the large Lipschitz constants of U_s . However, we have the weaker guarantee (whose proof follows that of Theorem 3) that performance is close to that predicted at the root.

Theorem 4 (FA Error with Weaker Bounds). If $L_{\infty}(\tilde{U}_s, \tilde{U}_s^{target}) \leq \epsilon$ for all $s \in S$, then $|R_1(\tilde{\pi}) - \widetilde{\mathsf{OPT}}_2| = \mathcal{O}(D\epsilon)$ where D is the depth of G and $\widetilde{\mathsf{OPT}}_2 = \max_{\mu_2} \tilde{U}_{root}(\mu_2)$.

Remark. The key technical difficulty here is finding \tilde{V} . In our experiments, $\tilde{\pi}_1^{\text{grim}}$ can be found analytically. In general large games, we can approximate $\tilde{\pi}_1^{\text{grim}}$, \tilde{V} by searching over S_2 , but use heuristics when expanding nodes in S_1 .

Implementation Details. (i) We use several techniques typically used to stabilize training such as target networks (Arulkumaran et al. 2017; Mnih et al. 2015) and prioritized experience replay (Schaul et al. 2015). (ii) In practice, instead of L_{∞} , we found it easier to train a loss based on the sum of the squared distances at the x-coordinate of the knots in \tilde{U}_s and \tilde{U}'_s , i.e., $L = \sum_{\mu_2 \in \{\text{knots}\}} [\tilde{U}_s(\mu_2) - \tilde{U}'_s(\mu_2)]^2$. Since L upper bounds L^2_{∞} , using it also avoids costly

promises and allows us to enjoys a similar FA guarantee. (iii) If G has a branching factor of β , then (1) in Algorithm 2 can be executed in $\mathcal{O}(\beta m)$ time. In practice, we use a brute force method better suited for batch GPU operations which runs in $\mathcal{O}((\beta m)^3)$. (iv) We train using only the decreasing portions of \tilde{U}_s . This does not lead to any loss in performance since payoffs in the increasing portion of an \tilde{U}_s are Pareto dominated. We do not want to 'waste' knots on learning the meaningless increasing portion. (v) Training trajectories were obtained by taking actions uniformly at random. Specifics for all implementation details are in the Appendix.

6 Experiments

We focused on the following two synthetic games. Game details and experiment environments are in the Appendix. Code is at https://github.com/lingchunkai/learn-epf-sefce.

Tantrum. TANTRUM is the game in Figure 2b repeated n times, with $q_1 > 0, q_2 \ge 1$, and rewards accumulated over stages. The only way P₁ can get positive payoffs is by threatening to throw a trantrum with the mutually destructive (-1, -1) outcome. Since $q_2 > 1$, P₂ has to use threats spanning over stages to sufficiently entice P₂ to accede. Even though TRANTRUM has $\mathcal{O}(3^n)$ leaves, it is clear that the grim (resp. altruistic) strategy is to throw (resp. not throw) a tantrum at every step. Hence \overline{V} and \underline{V} are known even when n is large, making TANTRUM a good testbed. The raw features f(s) is a 5-dimensional vector, the first 3 are the occurrences count of outcomes for previous stages, and the last 2 being a one-hot vector indicating the current state.

Resource Collection. RC is played on a $J \times J$ grid with a time horizon n. Each cell contains varying quantities of 2 different resources $r_1(x, y), r_2(x, y) \ge 0$, both of which are collected (at most once) by either players entering. Players begin in the center and alternately choose to either move to an adjacent cell or stay put. Each P_i is only interested in resource *i*, and players agree to pool together resources when the game ends. RC gives P_1 the opportunity to threaten P_2 with going 'on strike' if P_2 does not move to the cells that P_1 recommends. RC has approximately $\mathcal{O}(25^n)$ leaves. The grim strategy is for P_1 to stay put. However, unlike TANTRUM, computing <u>V</u> and \overline{V} still requires search (at least for P_2) at each state, which is still computationally expensive. We use as features (a) one-hot vector showing past visited locations, (b) the current coordinates of each player and whose turn it is (c) the amount of each resource collected, and (d) the number of rounds remaining.

6.1 Experimental Setup

Games with Fixed Parameters. We run 3 subexperiments. **[RC]** We experimented with RC with J = 7, n = 4 over 10 different games. Rewards r_i were generated using a log-Gaussian process over (x, y) to simulate spatial correlations (details in Appendix). We also report the payoffs from a 'non-strategic' P₁ which optimizes only for resources it collects, while letting P₂ best respond. **[TANTRUM]** We ran TANTRUM with $n = 25, q_1 = 1$ and q_2 chosen randomly. These games have > 1e12 states;



Figure 4: (a) Results for games with fixed parameters averaged over # specifies # trials. $\overline{\Delta_{OPT}}$, $\overline{\Delta_{SP}}$, and $\overline{\Delta_{non}}$ is the average difference between our method and the optimal SEFCE, subgame perfect Nash, and non-strategic leader commitment. (b)-(c) Learned EPFs at the root for RC. (d) A failure case in TANTRUM, even though learned policies are still near-optimal.

however, we can still obtain the optimal strategy due to the special structure of the game (note the subgame perfect equilibrium gives P₁ zero payoff). [**RC+**] We ran RC with J = 9, n = 6. Since G is large, we use approximates ($\tilde{\tau}, \tilde{V}, \tilde{V}$) obtained from $\tilde{\pi}_1^{\text{grim}}$ and $\tilde{\pi}^{\text{alt}}$. $\tilde{\pi}_1^{\text{grim}}$ is for P₂ to stay put, while \tilde{V} is obtained by applying search *online* (i.e., when *s* appears in training) for P₂ starting from *s*. Thus $\tilde{\tau}(s)$ can also be computed online from \tilde{V} . $\tilde{\pi}^{\text{alt}}$ is obtained by running exact search to a depth of 4 (counted from the root) and then switching to a greedy algorithm. On the rare occasion that $\tilde{V}(s) < \tilde{V}(s)$, we set $\tilde{V}(s) \leftarrow \tilde{V}(s)$. We report results in Figure 4a, which show the *difference* between P₁'s payoff for our method and (i) the optimal SEFCE, (ii) the subgame perfect Nash, and (iii) the non-strategic leader commitment.

Featurized TANTRUM. We allow q_1, q_2 to vary between stages of G, giving vectors $\mathbf{q}_i \in [1, \infty]^n$. Each trajectory uses different \mathbf{q}_i , which we append as features to our network, alongside the payoffs already collected for each player. For training, we draw i.i.d. samples of $\mathbf{q}_i^j \sim \exp(1) +$ 1. The evaluation metric is $\kappa = R_1(\tilde{\pi})/\text{OPT}$, i.e., the ratio of P₁'s payoffs under $\tilde{\pi}$ compared to the optimal π . For each *n*, we test on 100 **q**-vectors not seen during training and compare their κ against a 'greedy' strategy which recommends P₂ to accede as long as there are sufficient threats in the remainder of the game for P₁ (details in Appendix). We also stress test $\tilde{\pi}$ on a different *test* distribution $\hat{\mathbf{q}}_i^j \sim \exp(1) + 4$. We report results in Figure 5a and 5b.

6.2 Results and Discussion

For fixed parameter games whose optimal value can be computed, we observe near optimal performance which significantly outperforms other baselines. For [RC], the average value of each an improvement of .5 is approximately equal to moving an extra half move. In [TANTRUM], the subgame perfect equilibrium is vacuous as P_1 is unable to issue threats and gets a payoff of 0. In [RC+], we are unable to fully expand the game tree, however, we still significantly outperform the non-strategic baseline.

For featurized TANTRUM, we perform near-optimally for small n, even when stress tested with out-of-distribution **q**'s (Figure 5a). Performance drops as n becomes larger, which is natural as EPFs become more complex. While performance degrades as n increases, we still significantly outper-



Figure 5: Results for Featurized TRANTRUM as depth n varies, based on κ , the ratio of the leader's payoff to the true optimum. (a) grd- $\overline{\kappa}$ and str- $\overline{\kappa}$ denote results for the baseline greedy method and our results when stress tested with **q** drawn from a distribution from training. (b) Proportion of trials which give $\kappa < \kappa_{\text{thresh}}$.

form the greedy baseline. The stress test suggests that the network is not merely memorizing data.

Figures 4b and 4c shows the learned EPFs at the root for epochs 100k and 2M, obtained directly or from onestep lookahead. As explained in Section 5, we only learn the decreasing portions of EPFs. After 2M training epochs, the predicted EPFs and one-step lookahead mirrors the true EPF in the decreasing portions, which is not the case at the beginning. At the beginning of training, many knots (red markers) are wasted on learning the 'useless' increasing portions on the left. After 2M epochs, knots (blue markers) were learning the EPF at the 'useful' decreasing regions.

Figure 4d gives an state in TANTRUM whose EPF yields high loss even after training. This failure case is not rare since TANTRUM is large. Yet, the resultant action is still optimal—in this case the promise to P₂ was $\mu_2 = -25.5$ which is precisely $\overline{V}(s)$. Like MDPs, policies can be nearoptimal even with high Bellman losses in some states.

7 Conclusion

We proposed a novel method of performing FA on EPFs that allows us to efficiently solve for SEFCE. This is to the best of our knowledge, the first time a such an object has been learned from state features, leading to a FA-based method of solving Stackelberg games with performance guarantees. We hope that our approach will help to close the current gap between solving zero-sum and general-sum games.

Acknowledgements

This work was supported by NSF grant IIS-2046640 (CA-REER).

References

Arulkumaran, K.; Deisenroth, M. P.; Brundage, M.; and Bharath, A. A. 2017. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6): 26–38.

Bakhtin, A.; Wu, D.; Lerer, A.; and Brown, N. 2021. No-Press Diplomacy from Scratch. *Advances in Neural Information Processing Systems*, 34.

Bošanský, B.; Brânzei, S.; Hansen, K. A.; Lund, T. B.; and Miltersen, P. B. 2017. Computation of Stackelberg equilibria of finite sequential games. *ACM Transactions on Economics and Computation (TEAC)*, 5(4): 1–24.

Bosanský, B.; Brânzei, S.; Hansen, K. A.; Miltersen, P. B.; and Sørensen, T. B. 2015. Computation of Stackelberg Equilibria of Finite Sequential Games. *CoRR*, abs/1507.07677.

Bosansky, B.; and Cermak, J. 2015. Sequence-form algorithm for computing stackelberg equilibria in extensive-form games. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.

Brown, N.; and Sandholm, T. 2017. Libratus: the superhuman AI for no-limit poker. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*.

Brown, N.; and Sandholm, T. 2019. Superhuman AI for multiplayer poker. *Science*, 365(6456): 885–890.

Castelletti, A.; Pianosi, F.; and Restelli, M. 2011. Multiobjective Fitted Q-Iteration: Pareto frontier approximation in one single run. In 2011 International Conference on Networking, Sensing and Control, 260–265. IEEE.

Cermak, J.; Bosansky, B.; Durkota, K.; Lisy, V.; and Kiekintveld, C. 2016. Using correlated strategies for computing stackelberg equilibria in extensive-form games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.

Cermák, J.; Bošanský, B.; Durkota, K.; Lisý, V.; and Kiekintveld, C. 2016. Using Correlated Strategies for Computing Stackelberg Equilibria in Extensive-form Games. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, 439–445. AAAI Press.

Černý, J.; Bošanský, B.; and Kiekintveld, C. 2018. Incremental Strategy Generation for Stackelberg Equilibria in Extensive-Form Games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 151–168. ACM.

De Jonge, D.; and Zhang, D. 2020. Strategic negotiations for extensive-form games. *Autonomous Agents and Multi-Agent Systems*, 34(1): 1–41.

Dermed, M.; and Charles, L. 2013. Value methods for efficiently solving stochastic games of complete and incomplete information. Ph.D. thesis, Georgia Institute of Technology.

Dietterich, T.; and Wang, X. 2001. Batch value function approximation via support vectors. *Advances in neural information processing systems*, 14.

Fang, F.; Nguyen, T. H.; Pickles, R.; Lam, W. Y.; Clements, G. R.; An, B.; Singh, A.; Schwedock, B. C.; Tambe, M.; and Lemieux, A. 2017. PAWS-A Deployed Game-Theoretic Application to Combat Poaching. *AI Magazine*, 38(1): 23.

Foerster, J. N.; Chen, R. Y.; Al-Shedivat, M.; Whiteson, S.; Abbeel, P.; and Mordatch, I. 2017. Learning with opponentlearning awareness. *arXiv preprint arXiv:1709.04326*.

Gray, J.; Lerer, A.; Bakhtin, A.; and Brown, N. 2021. Human-Level Performance in No-Press Diplomacy via Equilibrium Search. In *International Conference on Learning Representations*.

Greenwald, A.; Hall, K.; Serrano, R.; et al. 2003. Correlated Q-learning. In *ICML*, volume 3, 242–249.

Karwowski, J.; and Mańdziuk, J. 2020. Double-oracle sampling method for Stackelberg Equilibrium approximation in general-sum extensive-form games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2054–2061.

Kasparov, G. 2018. Chess, a Drosophila of reasoning. *Science*, 362(6419): 1087–1087.

Lagoudakis, M. G.; and Parr, R. 2003. Least-squares policy iteration. *The Journal of Machine Learning Research*, 4: 1107–1149.

Lanctot, M.; Zambaldi, V.; Gruslys, A.; Lazaridou, A.; Tuyls, K.; Pérolat, J.; Silver, D.; and Graepel, T. 2017. A unified game-theoretic approach to multiagent reinforcement learning. *Advances in neural information processing systems*, 30.

Leibo, J. Z.; Zambaldi, V.; Lanctot, M.; Marecki, J.; and Graepel, T. 2017. Multi-agent reinforcement learning in sequential social dilemmas. *arXiv preprint arXiv:1702.03037*.

Letchford, J.; and Conitzer, V. 2010. Computing optimal strategies to commit to in extensive-form games. In *Proceedings of the 11th ACM conference on Electronic commerce*, 83–92. ACM.

Letchford, J.; MacDermed, L.; Conitzer, V.; Parr, R.; and Isbell, C. L. 2012. Computing optimal strategies to commit to in stochastic games. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*.

MacDermed, L.; Narayan, K. S.; Isbell, C. L.; and Weiss, L. 2011. Quick polytope approximation of all correlated equilibria in stochastic games. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*.

Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.

Moravčík, M.; Schmid, M.; Burch, N.; Lisỳ, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; and Bowling, M. 2017. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337): 508–513.

Munos, R.; and Szepesvári, C. 2008. Finite-Time Bounds for Fitted Value Iteration. *Journal of Machine Learning Research*, 9(5).

Murray, C.; and Gordon, G. 2007. *Finding correlated equilibria in general sum stochastic games.*

Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in PyTorch.

Perolat, J.; de Vylder, B.; Hennes, D.; Tarassov, E.; Strub, F.; de Boer, V.; Muller, P.; Connor, J. T.; Burch, N.; Anthony, T.; et al. 2022. Mastering the Game of Stratego with Model-Free Multiagent Reinforcement Learning. *arXiv preprint arXiv:2206.15378*.

Pérolat, J.; Strub, F.; Piot, B.; and Pietquin, O. 2017. Learning Nash equilibrium for general-sum Markov games from batch data. In *Artificial Intelligence and Statistics*, 232–241. PMLR.

Pita, J.; Jain, M.; Ordónez, F.; Portway, C.; Tambe, M.; Western, C.; Paruchuri, P.; and Kraus, S. 2008. ARMOR Security for Los Angeles International Airport. In *AAAI*, 1884–1885.

Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.

Schmid, M.; Moravcik, M.; Burch, N.; Kadlec, R.; Davidson, J.; Waugh, K.; Bard, N.; Timbers, F.; Lanctot, M.; Holland, Z.; et al. 2021. Player of games. *arXiv preprint arXiv:2112.03178*.

Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.

Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; et al. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419): 1140–1144.

Vinyals, O.; Babuschkin, I.; Chung, J.; Mathieu, M.; Jaderberg, M.; Czarnecki, W.; Dudzik, A.; Huang, A.; Georgiev, P.; Powell, R.; Ewalds, T.; Horgan, D.; Kroiss, M.; Danihelka, I.; Agapiou, J.; Oh, J.; Dalibard, V.; Choi, D.; Sifre, L.; Sulsky, Y.; Vezhnevets, S.; Molloy, J.; Cai, T.; Budden, D.; Paine, T.; Gulcehre, C.; Wang, Z.; Pfaff, T.; Pohlen, T.; Yogatama, D.; Cohen, J.; McKinney, K.; Smith, O.; Schaul, T.; Lillicrap, T.; Apps, C.; Kavukcuoglu, K.; Hassabis, D.; and Silver, D. 2019. AlphaStar: Mastering the Real-Time Strategy Game Star-Craft II. https://deepmind.com/blog/alphastar-masteringreal-time-strategy-game-starcraft-ii/. Accessed: 2022-02-01.

Wang, Y.; Shi, Z. R.; Yu, L.; Wu, Y.; Singh, R.; Joppa, L.; and Fang, F. 2019. Deep reinforcement learning for green security games with real-time information. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1401–1408.

Zhong, H.; Yang, Z.; Wang, Z.; and Jordan, M. I. 2021. Can Reinforcement Learning Find Stackelberg-Nash Equilibria in General-Sum Markov Games with Myopic Followers? *arXiv preprint arXiv:2112.13521*. Zinkevich, M.; Greenwald, A.; and Littman, M. 2005. Cyclic equilibria in Markov games. *Advances in neural information processing systems*, 18.

A Proof of Theorem 3

A.1 Preliminaries

The proof follows 2 steps. First, we show that the estimate \tilde{U}_s is not too different from the optimal U_s . Let the depth of a state D(s) be the longest path needed to reach a leaf, i.e.,

$$D(s) = \begin{cases} 0 & s \in \mathcal{L} \\ \max_{s' \in \mathcal{C}(s)} D(s') + 1 & \text{otherwise} \end{cases}$$

Since G is finite, $D = \max_{s \in S} D(s)$ and is finite.

Let \tilde{U}_s be our predicted EPF at state s. For $s \in S \setminus \mathcal{L}$, we denote (for this section only) using shorthand

$$\tilde{U}'_{s} = \tilde{U}^{\text{target}}_{s} = \begin{cases} \left[\bigwedge_{s' \in \mathcal{C}(s)} \tilde{U}_{s'} \right] (\mu) & \text{if } s \in \mathcal{S}_{1} \\ \left[\bigwedge_{s' \in \mathcal{C}(s)} \tilde{U}_{s'} \triangleright \tau(s') \right] (\mu) & \text{if } s \in \mathcal{S}_{2} \end{cases}$$

be what is obtained from one-step lookahead using Section 3, and for $s \in \mathcal{L}$,

$$\tilde{U}_s(\mu_2) = \tilde{U}'_s(\mu_2) = \begin{cases} r_1(s) & \mu_2 = r_2(s) \\ -\infty & \text{otherwise} \end{cases}.$$

For clarity, we denote likewise for the exact EPFs U'_s (which will be equal by definition to U_s).

Domains of EPFs Given any real valued function $h : \mathbb{R} \to \mathbb{R}$, we denote its domain by $Dom[h] = \{x | h(x) > -\infty\}$. The following lemmas ensure that the required *domains* all match. This theorem is added for completeness; the reader can skip over this section if desired.

Lemma 1. For all $s \in S$, $Dom[U_s] = Dom[U'_s] = Dom[\tilde{U}_s] = Dom[\tilde{U}'_s] = [\underline{V}(s), \overline{V}(s)]$,

The proofs are straightforward but trivial, hence they are deferred to Section C

Proof. The first equality was shown in Lemma 3. We can, in fact reuse the proof of Lemma 3 for \tilde{U}'_s and \tilde{U}_s . This completes this Lemma 1.

Hence, we no longer have to worry about mismatched domains.

Upper Concave Envelopes and Left Truncations First, observe that since \tilde{U}_s is a one-dimensional function, \bigwedge can be written alternatively as

$$\left[\bigwedge_{s'\in\mathcal{C}(s)} U_{s'}\right](\mu_2) = \max_{\substack{s',s''\in\mathcal{C}(s)\\t\in[0,1];\mu',\mu''\in\mathbb{R}\\t\mu'+(1-t)\mu''=\mu_2}} tU_{s'}(\mu') + (1-t)U_{s''}(\mu''),$$

that is, the maximum that could be obtained by interpolating between at most two points across 2 children states $s', s'' \in C(s)$ (this follows from the fact that all U's are 1-dimensional functions (Letchford and Conitzer 2010)).

Recall that $\tau(s')$ is defined only for $s' \in C(s)$, where $s \in S_2$. For convenience, we define $\beta(s') = \tau(s')$ when $s \in S_2$ and $-\infty$ when $s \in S_1$. β plays the same role as τ , since left truncating at $-\infty$ does not change anything, i.e., $f \triangleright (-\infty) = f$ for any $f : \mathbb{R} \to \mathbb{R}$. Using β allows us to perform a 'dummy' left truncation when $s \in S_1$ and avoid having to split into different cases.

Our Goal Our goal is to show that assuming that the L_{∞} loss is low for all states s, i.e.,

$$\max_{\mu_2 \in [\underline{V}(s), \overline{V}(s)]} |\tilde{U}_s(\mu_2) - \tilde{U}'_s(\mu_2)| \le \epsilon$$
(2)

then we will enjoy good performance, i.e, the leader gets a payoff of order $\mathcal{O}(\epsilon D)$ less than optimal (additively).

A.2 Learned EPFs are Approximately Optimal

The first half of the theorem is to show that our learned EPFs \tilde{U}_s are for all s, close to the true U_s pointwise. We prove the main theorem by strong induction on the states by increasing depth the following. and (1). Our induction hypothesis is

$$\mathcal{K}_j: \quad \max_{\mu_2 \in [\underline{V}(s), \overline{V}(s)]} |\tilde{U}_s(\mu_2) - U_s(\mu_2)| \leq j\epsilon \qquad \forall s \text{ where } D(s) = j.$$

By definition, K_0 satisfies our requirement since $s \in \mathcal{L}$. Thus, the base case is satisfied. Now we prove the inductive case. Assume that $\mathcal{K}_0, \dots, \mathcal{K}_{j-1}$ are all satisfied. We want to show \mathcal{K}_j using (2). **Lemma 2.** Let $s \in S$ such that D(s) = j. Suppose $\mathcal{K}_0, \ldots, \mathcal{K}_{j-1}$ are true. Then we have $|\tilde{U}'_s(\mu_2) - U'_s(\mu_2)| \le \epsilon(j-1)$ for all $\mu_2 \in [\underline{V}(s), \overline{V}(s)].$

Proof. Fix μ_2 . We want to show $|\tilde{U}'_s(\mu_2) - U'_s(\mu_2)| \leq \epsilon(j-1)$. Now, let $\hat{\sigma} = (\hat{s}', \hat{s}'', \hat{t}, \hat{\mu}', \hat{\mu}'')$ be the parameters which achieves the maximum

$$\underset{\substack{s',s'' \in \mathcal{C}(s) \\ t \in [0,1]; \mu', \mu'' \in \mathbb{R} \\ t\mu' + (1-t)\mu'' = \mu_2}}{ \arg\max} t[U_{s'} \triangleright \beta(s')](\mu') + (1-t)[U_{s''} \triangleright \beta(s'')](\mu'')$$
(3)

and similarly when we are working with learned EFPs, $\tilde{\sigma} = (\tilde{s}', \tilde{s}'', \tilde{t}, \tilde{\mu}', \tilde{\mu}'')$

$$\underset{\substack{s',s'' \in \mathcal{C}(s) \\ t \in [0,1]; \mu', \mu'' \in \mathbb{R} \\ t\mu' + (1-t)\mu'' = \mu_2 }}{ \arg\max} t[\tilde{U}_{s'} \triangleright \beta(s')](\mu') + (1-t)[\tilde{U}_{s''} \triangleright \beta(s'')](\mu''),$$
(4)

which are the arguments that optimized give $U'_s(\mu_2)$ and $\tilde{U}'_s(\mu_2)$ respectively. That is, $\hat{\sigma}$ and $\tilde{\sigma}$ gives how the point at the EPF with the x coordinate equal to μ_2 is obtained as a mixture of at most 2 points from the upper convex envelope (when s' = s'', we simply repeat the 2 points and set t = 1/2 for simplicity). We proceed by showing a contradiction. We have two cases.

Case 1. Suppose that $\tilde{U}'_s(\mu_2) > U'_s(\mu_2) + \epsilon(j-1)$. Then we have

$$\left| \underbrace{\tilde{t}[\tilde{U}_{\tilde{s}'} \triangleright \beta(\tilde{s}')](\tilde{\mu}') + (1 - \tilde{t})[\tilde{U}_{\tilde{s}''} \triangleright \beta(\tilde{s}'')](\tilde{\mu}'')}_{=\tilde{U}'_{s}(\mu_{2}) \geq U'_{s}(\mu_{2}) + \epsilon(j-1)} - \underbrace{\tilde{t}[U_{\tilde{s}'} \triangleright \beta(\tilde{s}')](\tilde{\mu}') + (1 - \tilde{t})[U_{\tilde{s}''} \triangleright \beta(\tilde{s}'')](\tilde{\mu}'')}_{\leq U'_{s}(\mu_{2})} \right|$$

$$= \left| \tilde{t} \left([\tilde{U}_{\tilde{s}'} \triangleright \beta(\tilde{s}')](\tilde{\mu}') - [U_{\tilde{s}'} \triangleright \beta(\tilde{s}')](\tilde{\mu}') \right) + (1 - \tilde{t}) \left([\tilde{U}_{\tilde{s}''} \triangleright \beta(\tilde{s}'')](\tilde{\mu}'') - [U_{\tilde{s}''} \triangleright \beta(\tilde{s}'')](\tilde{\mu}'') \right) \right|$$

$$\leq \epsilon(j-1)$$
(5)

and where first inequality inside $|\cdot|$ follows from our assumption in case 1, and the second from the fact that $U'_{s}(\mu_{2})$ was taken from an argmax, i.e., (3). The third line holds from our induction hypothesis \mathcal{K}_i (2), where $i \in [0, j-1]$, the fact that D(s') < D(s) = j and how the \triangleright operator cannot increase the absolute error of the difference.⁷ However, these 3 inequalities cannot hold simultaneously, thus the assumption for Case 1 cannot be true, i.e., we must have $\tilde{U}'_s(\mu_2) \le U'_s(\mu_2) + \epsilon(j-1)$

Case 2. Suppose that $\tilde{U}'_s(\mu_2) < U'_s(\mu_2) - \epsilon(j-1)$. Then using a similar derivation we have

$$\frac{\left| \hat{t}[U_{\hat{s}'} \triangleright \beta(\hat{s}')](\hat{\mu}') + (1-\hat{t})[]U_{\hat{s}''} \triangleright \beta(\hat{s}'')](\hat{\mu}'')}{=U'_{s}(\mu_{2}) > \tilde{U}'_{s}(\mu_{2}) + \epsilon(j-1)} - \hat{t}[\tilde{U}_{\hat{s}'} \triangleright \beta(\hat{s}')](\hat{\mu}') + (1-\hat{t})[\tilde{U}_{\hat{s}''} \triangleright \beta(\hat{s}'')](\hat{\mu}'')}{\leq \tilde{U}'_{s}(\mu_{2})} \right|$$

$$= \left| \hat{t} \left([U_{\hat{s}'} \triangleright \beta(\hat{s}')](\hat{\mu}') - [\tilde{U}_{\hat{s}'} \triangleright \beta(\hat{s}')](\hat{\mu}') \right) + (1-\hat{t}) \left([U_{\hat{s}''} \triangleright \beta(\hat{s}'')](\hat{\mu}'') - [\tilde{U}_{\hat{s}''} \triangleright \beta(\hat{s}'')](\hat{\mu}'') \right) \right|$$

$$\leq \epsilon(j-1)$$
(6)

where the first inequality inside $|\cdot|$ comes from case 2's assumption, the second is from the fact that $\tilde{U}'_s(\mu_2)$ was taken from an argmax, i.e., (4). The third line follows from the induction hypothesis \mathcal{K}_i , where $i \in [0, j - 1]$ and the depth of s. These 3 inequalities cannot hold simultaneously, so our assumption in case 2 cannot be true and $\tilde{U}'_s(\mu_2) \geq U'_s(\mu_2) - \epsilon(j-1)$.

Combining the result from both cases gives the desired result.

⁷Note that since $\tilde{\sigma}$ is the argmax, we are guaranteed to not have values of $\tilde{\mu}$ that lie outside the domains of the truncated EPFs.

Now, we consider the μ_2 which has the worst possible discrepancy, which gives

$$\max_{\substack{\mu_{2} \in [\underline{V}(s), \overline{V}(s)]}} |U_{s}(\mu_{2}) - U_{s}(\mu_{2})| \\
\leq \max_{\substack{\mu_{2} \in [\underline{V}(s), \overline{V}(s)]}} |\tilde{U}_{s}(\mu_{2}) - \tilde{U}_{s}'(\mu_{2})| + |\tilde{U}_{s}'(\mu_{2}) - U_{s}(\mu_{2})| \\
= \max_{\substack{\mu_{2} \in [\underline{V}(s), \overline{V}(s)]}} |\tilde{U}_{s}(\mu_{2}) - U_{s}'(\mu_{2})| + |\tilde{U}_{s}'(\mu_{2}) - U_{s}'(\mu_{2})| \\
\leq \max_{\substack{\mu_{2} \in [\underline{V}(s), \overline{V}(s)]}} |\tilde{U}_{s}(\mu_{2}) - \tilde{U}_{s}'(\mu_{2})| + \max_{\substack{\mu_{2} \in [\underline{V}(s), \overline{V}(s)]}} |\tilde{U}_{s}'(\mu_{2}) - U_{s}'(\mu_{2})| \\
\leq \max_{\substack{\mu_{2} \in [\underline{V}(s), \overline{V}(s)]}} |\tilde{U}_{s}'(\mu_{2}) - U_{s}'(\mu_{2})| + \epsilon \\
\leq \epsilon j,$$
(7)

where the second line follows from the triangle inequality, the third line using the equality between $U_s(\mu_2)$ and $U'_s(\mu_2)$, the fourth from the fact that $\max_x |f(x) + g(x)| \le \max_x |f(x)| + \max_y |g(y)|$, the fifth from the assumption (2), the the last line from Lemma 2. The main theorem follows by induction on D(s), the fact that D(s) is bounded by the depth of the tree, and the fact that the base case \mathcal{K}_0 is trivially true.

Theorem 5. If $L_{\infty}(\tilde{U}_s, \tilde{U}'_s) \leq \epsilon$ for all $s \in S$, then $\max_{\mu_2 \in [\underline{V}(s), \overline{V}(s)]} |\tilde{U}_s(\mu_2) - U_s(\mu_2)| \leq \epsilon D$, where D is the depth of the game.

Proof. This follows by the definition of L_{∞} and the above derivations.

A.3 Leader's Payoffs from Induced Strategy is Close-To-Expected

Theorem 5 tells us that our EPF everywhere is close (pointwise) to the true EPF if ϵ is small. Now we need to establish the suboptimality when *playing* according to $\tilde{\pi}$, which is a joint policy implicit from \tilde{U}_s . We begin with some notation.

Let $\tilde{Q}_s(\mu_2) : [\underline{V}(s), \overline{V}(s)] \mapsto \mathbb{R}$ be the payoff to P_1 assuming we started at state s, promised a payoff of μ_2 to P_2 and used the approximate EPFs \tilde{U}_s for all descendent states $s' \supseteq s$ (the domain precludes unfulfilled promises). That is, given $\tilde{\sigma} = (\tilde{s}', \tilde{s}'', \tilde{t}, \tilde{\mu}', \tilde{\mu}'')$ given by

$$\underset{\substack{s',s'' \in \mathcal{C}(s) \\ t \in [0,1]; \mu', \mu'' \in \mathbb{R} \\ t\mu' + (1-t)\mu'' = \mu_2 }}{ \arg\max} t[\tilde{U}_{s'} \triangleright \beta(s')](\mu') + (1-t)[\tilde{U}_{s''} \triangleright \beta(s'')](\mu''),$$
(8)

the induced policy is to play to \tilde{s}' with probability \tilde{t} and a consequent promise of $\tilde{\mu}'$, as well as playing to \tilde{s}'' with probability $(1 - \tilde{t})$ and a promised payoff of $\tilde{\mu}''$. By definition, if $s \in \mathcal{L}$, $\tilde{Q}_s = \tilde{U}_s = U_s$ trivially. We also have the following recursive equations for a given $s \notin \mathcal{L}$, μ_2

$$\tilde{Q}_{s}(\mu_{2}) = \tilde{t}\tilde{Q}_{\tilde{s}'}(\tilde{\mu}') + (1 - \tilde{t})\tilde{Q}_{\tilde{s}''}(\tilde{\mu}'').$$
(9)

Theorem 6. $\left| \tilde{Q}_s(\mu_2) - \tilde{U}_s(\mu_2) \right| \le \epsilon D$ for all $s \in S$ and for all $\mu_2 \in [\underline{V}(s), \overline{V}(s)]$.

Proof. The proof is given by strong induction on the D(s) again. Let

$$\mathcal{H}_j: \quad \left| \tilde{Q}_s(\mu_2) - \tilde{U}_s(\mu_2) \right| \le \epsilon j \quad \forall s \text{ where } D(s) = j \tag{10}$$

By definition \mathcal{H}_0 is true. Now let us suppose that $\mathcal{H}_0 \dots \mathcal{H}_{j-1}$ is true. We have for $s \in \mathcal{S}, D(s) = j$,

$$\begin{split} & \left| \tilde{Q}_{s}(\mu_{2}) - \tilde{U}_{s}(\mu_{2}) \right| \\ \leq \left| \tilde{Q}_{s}(\mu_{2}) - \tilde{U}'_{s}(\mu_{2}) \right| + \left| \tilde{U}'_{s}(\mu_{2}) - \tilde{U}_{s}(\mu_{2}) \right| \\ \leq \left| \tilde{Q}_{s}(\mu_{2}) - \tilde{U}'_{s}(\mu_{2}) \right| + \epsilon \\ = \left| \tilde{t} \tilde{Q}_{\tilde{s}'}(\tilde{\mu}') + (1 - \tilde{t}) \tilde{Q}_{\tilde{s}''}(\tilde{\mu}'') - \tilde{t} [\tilde{U}_{\tilde{s}'} \triangleright \beta(\tilde{s}')](\tilde{\mu}') - (1 - \tilde{t}) [\tilde{U}_{\tilde{s}''} \triangleright \beta(\tilde{s}'')](\tilde{\mu}'') \right| + \epsilon \\ \leq \tilde{t} \underbrace{\left| \tilde{Q}_{\tilde{s}'}(\tilde{\mu}') - \tilde{U}_{\tilde{s}'}(\tilde{\mu}') \right|}_{\leq \epsilon(j-1)} + (1 - \tilde{t}) \underbrace{\left| \tilde{Q}_{\tilde{s}''}(\tilde{\mu}'') - \tilde{U}_{\tilde{s}''}(\tilde{\mu}'') \right|}_{\leq \epsilon(j-1)} + \epsilon \\ \leq \epsilon_{j} \end{split}$$

The second line follows from the triangle inequality. The third line follows from our FA assumption (2). The fourth line follows from expansion of the definitions of \tilde{Q}_s and \tilde{U}'_s , i.e., (9) and (4) The fifth line follows the induction hypothesis and the fact that $s', s'' \in \mathcal{C}(s)$ have at least one lower depth than s. Also, the truncation operator never causes any element to exceed domain bounds (which would give $-\infty$ values). By strong induction \mathcal{H}_j is true for all $j \in [0, D]$ and the theorem follows through directly.

A.4 Piecing Everything Together

Let $\mu_2^* = \arg \max_{\mu_2} U_{\text{root}}(\mu_2)$, i.e., the promise given to the follower *at the root* under the optimal policy π^* . Let $\tilde{\mu}_2 = \arg \max_{\mu_2} \tilde{U}_{\text{root}}(\mu_2)$, which is the promise to be given to the follower *at the root* under $\tilde{\pi}$. We have

$$U_{\text{root}}(\mu_{2}^{*}) - \tilde{Q}_{\text{root}}(\tilde{\mu}_{2})$$

$$= \underbrace{U_{\text{root}}(\mu_{2}^{*}) - \tilde{U}_{\text{root}}(\mu_{2}^{*})}_{|\cdot| \leq \epsilon D} + \underbrace{\tilde{U}_{\text{root}}(\mu_{2}^{*})}_{\leq \tilde{U}_{\text{root}}(\tilde{\mu}_{2})} - \tilde{Q}_{\text{root}}(\tilde{\mu}_{2})$$

$$\leq \epsilon D + \left| \tilde{U}_{\text{root}}(\tilde{\mu}_{2}) - \tilde{Q}_{\text{root}}(\tilde{\mu}_{2}) \right|$$

$$\leq 2\epsilon D.$$

The inequalities in the second line come from Theorem 5 and the definition of $\tilde{\mu}_2$; specifically that it is taken over the argmax. The last line comes from Theorem 6. To complete the proof, we simply observe that $\tilde{Q}_{root}(\tilde{\mu}_2)$ is precisely $r_1(\tilde{\pi})$ by definition.

B Proof of Theorem 4

The proof is essentially the same as presented in Section A.3, except that we are working with approximate bounds rather than strict ones. We have, using the new bounds,

For $s \in \mathcal{S} \setminus \mathcal{L}$, we denote using shorthand

$$\tilde{U}'_{s} = \tilde{U}^{\text{target}}_{s} = \begin{cases} \left[\bigwedge_{s' \in \mathcal{C}(s)} \tilde{U}_{s'} \right] (\mu) & \text{if } s \in \mathcal{S}_{1} \\ \left[\bigwedge_{s' \in \mathcal{C}(s)} \tilde{U}_{s'} \triangleright \tilde{\tau}(s') \right] (\mu) & \text{if } s \in \mathcal{S}_{2} \end{cases}$$

be what is obtained from one-step lookahead using Section 3, and for $s \in \mathcal{L}$,

$$\tilde{U}_s(\mu_2) = \tilde{U}'_s(\mu_2) = \begin{cases} r_1(s) & \mu_2 = r_2(s) \\ -\infty & \text{otherwise} \end{cases}.$$

This is the same as the case with exact $\underline{V}_s, \overline{V}_s$, but with a stricter truncation $\tilde{\tau}(s')$ for each $s' \in C(s)$. We define $\tilde{\beta}$ just like before: $\tilde{\beta}(s') = \tilde{\tau}(s')$ when $s \in S_2$ and $-\infty$ when $s \in S_1$. We follow along the same way as Theorem 3 in Section A.3.

Let $\tilde{Q}_s(\mu_2) : [V(s), \tilde{V}(s)] \mapsto \mathbb{R}$ be the payoff to P_1 assuming we started at state s, promised a payoff of μ_2 to P_2 and used the approximate EPFs \tilde{U}_s for all descendent states $s' \supseteq s$ (the domain precludes unfulfilled promises). That is, given $\tilde{\sigma} = (\tilde{s}', \tilde{s}'', \tilde{t}, \tilde{\mu}', \tilde{\mu}'')$ given by

$$\underset{\substack{s',s'' \in \mathcal{C}(s) \\ t \in [0,1]; \mu', \mu'' \in \mathbb{R} \\ t\mu' + (1-t)\mu'' = \mu_2}}{ \arg\max} t[\tilde{U}_{s'} \triangleright \tilde{\beta}(s')](\mu') + (1-t)[\tilde{U}_{s''} \triangleright \tilde{\beta}(s'')](\mu''),$$
(11)

the induced policy is to play to \tilde{s}' with probability \tilde{t} and a consequent promise of $\tilde{\mu}'$, as well as playing to \tilde{s}'' with probability $(1 - \tilde{t})$ and a promised payoff of $\tilde{\mu}''$. By definition, if $s \in \mathcal{L}$, $\tilde{Q}_s = \tilde{U}_s = U_s$ trivially. Just like before, we also have the following recursive equations for a given $s \notin \mathcal{L}$, μ_2

$$\tilde{Q}_{s}(\mu_{2}) = \tilde{t}\tilde{Q}_{\tilde{s}'}(\tilde{\mu}') + (1-\tilde{t})\tilde{Q}_{\tilde{s}''}(\tilde{\mu}'').$$
(12)

Let our induction hypothesis be

$$\mathcal{H}_j: \quad \left| \tilde{Q}_s(\mu_2) - \tilde{U}_s(\mu_2) \right| \le \epsilon j \quad \forall s \text{ where } D(s) = j.$$
(13)

By definition \mathcal{H}_0 is true. Now let us suppose that $\mathcal{H}_0 \dots \mathcal{H}_{j-1}$ is true. We have for $s \in \mathcal{S}, D(s) = j$,

$$\begin{split} & \left| \tilde{Q}_{s}(\mu_{2}) - \tilde{U}_{s}(\mu_{2}) \right| \\ \leq \left| \tilde{Q}_{s}(\mu_{2}) - \tilde{U}_{s}'(\mu_{2}) \right| + \left| \tilde{U}_{s}'(\mu_{2}) - \tilde{U}_{s}(\mu_{2}) \right| \\ \leq \left| \tilde{Q}_{s}(\mu_{2}) - \tilde{U}_{s}'(\mu_{2}) \right| + \epsilon \\ = \left| \tilde{t} \tilde{Q}_{\tilde{s}'}(\tilde{\mu}') + (1 - \tilde{t}) \tilde{Q}_{\tilde{s}''}(\tilde{\mu}'') - \tilde{t} [\tilde{U}_{\tilde{s}'} \triangleright \tilde{\beta}(\tilde{s}')](\tilde{\mu}') - (1 - \tilde{t}) [\tilde{U}_{\tilde{s}''} \triangleright \tilde{\beta}(\tilde{s}'')](\tilde{\mu}'') \right| + \epsilon \\ \leq \tilde{t} \underbrace{\left| \tilde{Q}_{\tilde{s}'}(\tilde{\mu}') - \tilde{U}_{\tilde{s}'}(\tilde{\mu}') \right|}_{\leq \epsilon(j-1)} + (1 - \tilde{t}) \underbrace{\left| \tilde{Q}_{\tilde{s}''}(\tilde{\mu}'') - \tilde{U}_{\tilde{s}''}(\tilde{\mu}'') \right|}_{\leq \epsilon(j-1)} + \epsilon \\ \leq \epsilon_{ij} \end{split}$$

The second line follows from the triangle inequality. The third line follows from our FA assumption (2). The fourth line follows from expansion of the definitions of \tilde{Q}_s and \tilde{U}'_s , i.e., (9) and (4) The fifth line follows the induction hypothesis and the fact that $s', s'' \in \mathcal{C}(s)$ have at least one lower depth than s. Also, the truncation operator never causes any element to exceed domain bounds (which would give $-\infty$ values). By strong induction \mathcal{H}_j is true for all $j \in [0, D]$. Finally, we observe that $\tilde{Q}_s(\mu_2) = R_1(\tilde{\pi})$ when $\tilde{\mu}_2 = \arg \max_{\mu_2} \tilde{U}_s(\mu_2)$. This completes the proof.

C Proof of Lemma 1

Lemma 3. For all $s \in S$, $Dom[U_s] = Dom[U's] = [\underline{V}(s), \overline{V}(s)]$.

Proof. The first equality is by definition. We now show that $Dom[U_s] = [\underline{V}, \overline{V}]$ by definition. Consider the state s we are applying (1) to and the 2 possible cases.

Case 1: $s \in S_1$. By definition $\overline{V}(s) = \max_{s' \in \mathcal{C}(s)} \overline{V}(s')$, and $\underline{V}(s) = \min_{s' \in \mathcal{C}(s)} \underline{V}(s)$. First, observe that

$$\max \left\{ Dom \left[U'_{s} \right] \right\} = \max \left\{ Dom \left[\bigwedge_{s' \in \mathcal{C}(s)} U_{s'} \right] \right\}$$
$$= \max_{s' \in \mathcal{C}(s)} \max \left\{ Dom \left[U_{s'} \right] \right\}$$
$$= \max_{s' \in \mathcal{C}(s)} \overline{V}(s')$$
$$= \overline{V}(s),$$
(14)

where the second line follows from the fact that the largest x-coordinate after taking the upper-concave-envelope is the largest of the largest-x coordinates over each $U_{s'}$. Similarly, we have

$$\min \left\{ Dom \left[U'_{s} \right] \right\} = \min \left\{ Dom \left[\bigwedge_{s' \in \mathcal{C}(s)} U_{s'} \right] \right\}$$
$$= \min_{s' \in \mathcal{C}(s)} \min \left\{ Dom \left[U_{s'} \right] \right\}$$
$$= \min_{s' \in \mathcal{C}(s)} \underline{V}(s')$$
$$= \underline{V}(s),$$
(15)

where the second line comes again from the fact that the lowest x-coordinate after taking upper concave envelopes is the smallest of all the smallest x-coordinates over each $U_{s'}$. Now, (14) and (15) established the lower and upper limits of U'_s . Since U'_s is concave, for every $\min \{Dom[U'_s]\} \leq \mu_2 \leq \max \{Dom[U'_s]\}\)$ we have $U'_s(\mu_2) \geq \min (U'_s(\min \{Dom[U'_s]\}), U'_s(\max \{Dom[U'_s]\})) > -\infty$. This completes Case 1.

Case 2: $s \in S_2$. By definition, $\overline{V}(s) = \max_{s' \in C(s)} \overline{V}(s')$ and $\underline{V}(s) = \max_{s' \in C(s)} \underline{V}(s')$ (note the difference with Case 1, since P₂ decides the current action). We again work out upper and lower bounds of $Dom[U'_s]$.

$$\max \left\{ Dom \left[U'_{s} \right] \right\} = \max \left\{ Dom \left[\bigwedge_{s' \in \mathcal{C}(s)} U_{s'} \triangleright \tau(s') \right] \right\}$$
$$= \max_{s' \in \mathcal{C}(s)} \max \left\{ Dom \left[U_{s'} \triangleright \tau(s') \right] \right\}$$
$$= \max_{s' \in \mathcal{C}(s)} \max \left\{ \left[\underline{V}(s'), \overline{V}(s') \right] \cap [\tau(s'), \infty) \right\}$$
$$= \max_{s' \in \mathcal{C}(s)} \overline{V}(s')$$
$$= \overline{V}(s),$$
(16)

where the fourth line follows from the fact that $= \max_{s' \in \mathcal{C}(s)} \overline{V}(s') \ge \max_{s^{!} \in \mathcal{C}(s); s^{!} \neq s'} \overline{V}(s^{!}) \ge \max_{s^{!} \in \mathcal{C}(s); s^{!} \neq s'} \underline{V}(s^{!}) = \tau(s')$ (i.e., that the highest x coordinate in U'_{s} is never part of the left-truncation step). For any $s' \in \mathcal{C}(s)$,

$$Dom\left[U_{s'} \triangleright \tau(s')\right] = \begin{cases} \emptyset & \max\left\{Dom[U_{s'}]\right\} < \tau(s') \\ Dom\left[U_{s'}\right] & \tau(s') < \min\left\{Dom\left[U_{s'}\right]\right\} \\ [\tau(s'), \max\left\{Dom[U_{s'}]\right\}\right] & \text{otherwise} \end{cases}$$
(17)

For s' where $Dom[U_{s'} \triangleright \tau(s')] \neq \emptyset$, we have

$$\min \left\{ Dom \left[U_{s'} \triangleright \tau(s') \right] \right\} = \max_{s'' \in \mathcal{C}(s)} \min \left\{ Dom \left[U_{s''} \right] \right\}.$$
(18)

Note that this is not dependent on s'. Also, note that it cannot be the case that $Dom[U_{s'} \triangleright \tau(s')] = \emptyset$ for all s'. In particular, consider $s^* = \arg \max_{s' \in \mathcal{C}(s)} \max \{Dom[U_{s'}]\}$, clearly, $\max \{Dom[U_{s'}]\} \ge \tau(s^*)$ so we do not wind up with the empty set in (17). For simplicity, let $\min \emptyset = \infty$. Hence, we can write

$$\min \left\{ Dom \left[U'_{s} \triangleright \tau(s') \right] \right\} = \min \left\{ Dom \left[\bigwedge_{s' \in \mathcal{C}(s)} U_{s'} \triangleright \tau(s') \right] \right\}$$
$$= \min_{s' \in \mathcal{C}(s)} \min \left\{ Dom \left[U_{s'} \triangleright \tau(s') \right] \right\}$$
$$= \min_{s' \in \mathcal{C}(s)} \max_{s'' \in \mathcal{C}(s)} \min \left\{ Dom \left[U_{s''} \right] \right\}$$
$$= \max_{s'' \in \mathcal{C}(s)} \min \left\{ Dom \left[U_{s''} \right] \right\}$$
$$= \underbrace{V}(s).$$
(19)

The first line is by definition. The second line uses the same argument as in case 1. The second line follows (18) and the fact that at least one s^* exists. The last line follows from the definition of $\underline{V}(s)$. As with case 1, we use (16), (19) and the fact that U'_s is concave to show that $U'_s(\mu_2) > -\infty$ for $\mu_2 \in [\underline{V}(s), \overline{V}(s)]$. This completes the proof.

We are now ready to tackle the proof of Lemma 1 (reproduced here): For all $s \in S$,

$$Dom[U_s] = Dom[U'_s] = Dom[\tilde{U}_s] = Dom[\tilde{U}'_s] = [\underline{V}(s), \overline{V}(s)],$$

Proof. The first equality was shown in Lemma 3. We can, in fact reuse the proof of Lemma 3 by replacing U'_s and U_s with \tilde{U}'_s and \tilde{U}_s . This completes this Lemma 1.

D Extensions to Games with Chance

For games with chance, backups will involve infimal convolutions (Cermák et al. 2016). Denote the set of chance nodes by S_C . For $s \in S_C$ and we denote the probability of transition for from s to s' to be $\pi_C(s', s)$. The set of equations at (1) has to be augmented by the case where $s \in S_C$. In these cases, we have

$$U_s(\mu) = \bigoplus_{s' \in \mathcal{C}(s)} \pi_{\mathsf{C}}(s', s) U_{s'}(\mu/\pi_{\mathsf{C}}(s', s)),$$



Figure 7: Sample game for difference between SSE and SEFCE.



Figure 8: From left to right: EPFs based on the game in Figure 6. (a) Enforceable EPF at s'', (b) EPF for SSE at s, (c) EPF for SEFCE at s. The EPF at s' is the same as that in Figure 1b.

where \bigoplus is the maximal-convolution operator (similar to the inf-conv operator for convex functions),

$$f_1 \bigoplus f_2(\mu) = \sup_{y} \left\{ f_1(\mu - y) + f_2(y) | y \in \mathbb{R} \right\}.$$

Refer to (Cermák et al. 2016) for more details as to why \bigoplus is the right operator to be used. It is well known that \bigoplus can be efficiently implemented (linear in the number of knots) when functions are piecewise linear concave (sort the line segments in all f_i based on gradients and stitch these line segments together in ascending order of gradients). Thankfully, this holds for EPFs. Furthermore, applying \bigoplus to piecewise linear concave functions gives another piecewise linear concave function. Hence, EPFs of \tilde{U}_s for each state and trained again using the L_{∞} loss. Theorem 3 still holds with some minor additions (we omit the proof in this paper).

E Comparison between EPFs between SSE and SEFCE

One of the key disadvantages of EPFs in SSE is that they could be non-concave, or worst still, discontinuous. Consider the game in Figure 6 and the EPFs in Figure 8. In Figure 8b, we can see that the EPF is neither concave or even continuous. This is because the SSE takes pointwise maximums at follower nodes and not upper-concave envelopes. On the other hand, Figure 8c shows the EPF of an SEFCE. Here, it is much better behaved, being a piecewise linear concave function.Furthermore, as mentioned in Theorem 1, the EPF in SEFCE dominates (is always higher or equal to at all x-coordinates) the EPF of SSE. This means the SEFCE can give P_1 more payoff than SSE. Also, every SSE is an SEFCE, but no vice versa.

See (Bošanskỳ et al. 2017) for a breakdown of computational complexity for different classes of games, (e.g., correlation/signaling (correlated, pure, behavioral), whether there is chance, and different levels of imperfect information).

F A Useful Way of Reasoning about SEFCE

For readers who are more familiar with SSE or are uncomfortable with the 'correlation' present in SEFCE, we give an easy interpretation of SEFCE in perfect information games. Given G, consider a modified game G', where before a follower vertex s, we add a leader vertex s' just before it with *two* actions, where each action leads to a copy of the game rooted at the follower vertex s. This construction is repeatedly performed in a bottom-up fashion. After this entire process is completed, we will find the *SSE* of G' (which is a much larger game than G).

The only purpose of this leader vertex is to allow mixing between follower strategies. Now, the recommendation to the follower is explicit via s'. Each action in s' corresponds to a single recommended action for the follower at s. Note that only a binary signal is needed (since mixing will only occur between at most 2 follower actions). Let s_a and s_b be duplicate follower vertices. Crucially, the leader is allowed to commit to different strategies for each subgame following s_a and s_b . The SSE in G' can be mapped to the SEFCE in G. Since in SSE, best responses are pure, and the probabilities leading to s_a and s_b (from the leader signaling node) give the probabilities at state s in G for the (pure) action to be taken at s_a and s_b .



Figure 9: Example of duplicated follower vertices based off the game in Figure 6.

We reiterate that this construction is not one that is practical, but rather one to help to gain intuition for the SEFCE. The solution using the constructed game is the same as the SEFCE by running the algorithm in Section 3 and seeing how the \bigwedge operator (for SSE) at the added root in G' mimics the follower vertex in SEFCE.

G Implementation Details

G.1 Borrowing Techniques from RL and FVI

We employ target networks (Mnih et al. 2015). Instead of performing gradient descent on the 'true' loss, we create a 'frozen' copy of the network which we use the compute \tilde{U}'_s (i.e., $\tilde{U}^{\text{target}}_s$). However, \tilde{U} is still computed from the main network (with weights to be updated in gradient descent). The key idea is that the target \tilde{U}'_s is no longer update at every epoch, which can destabilize training. We update the target network with the main network once every 2000 episodes.

For larger games, we noticed that the bulk of loss was attributed to a small fraction of states. To focus attention on these states, we employ prioritized replay (Schaul et al. 2015). We set the probability of selecting each state s in the replay buffer to be proportionate to the square root of the last loss, i.e., $L(U_s, U'_s)^{\alpha}$ observed. We used $\alpha = 0.5$ for convenience ((Schaul et al. 2015) suggest a value of 0.7).

Finding out the best hyperparameters for these add-ons is beyond the scope of this paper and left as future work.

G.2 Modified Loss Function

Our experience is that L_{∞} does manage to learn EPFs well, however, learning can be slow and sometimes unstable. Our hypothesis is that slow learning is due to the fact only the point responsible for the loss, as well as its neighbors has its coordinates updated during training. This very 'local' learning of \tilde{U}_s makes learning slow, particularly at the start of training. Second, we found that rather than L_{∞} , using the *square* of the largest absolute pointwise difference tends to stabilize training (though Theorem 3 would have to be modified to be in terms of $\sqrt{\epsilon}$ instead).

Let U_s and U'_s be two EPFs represented by k_1 and k_2 knots. Let $X_1 = \{x_1, \dots, x_{k_1}\}$ and $X_2 = \{x'_1, \dots, x'_{k_2}\}$ be the x-coordinates of the knots in U_s and U'_s respectively. Then, we use the following loss

$$L(U_s, U'_s) = \sum_{x \in X_1 \cup X_2} \left(U_s(x) - U'_s(x) \right)^2.$$
(20)

This loss still avoids costly promises since we are still taking pointwise differences (rather than over an integral).

G.3 Practical Implementation of Upper Concave Envelope and Left-Truncation

Theoretically, finding the upper concave envelope of k points can be found in linear time. However, this algorithm requires a significant number of backtracking and if-else statements, making this implementation unsuitable for batch operations on a GPU. The alternative which we employ runs in $\mathcal{O}(k^3)$ time which is in practice much faster when run on a GPU. For every distinct pairs of points (x_i, y_i) , and (x_j, y_j) , we check, for every point (x_a, y_a) where $x_a \in [x_i, x_j]$ whether (x_a, y_a) lies below or above the line segment $(x_i, y_i), (x_j, y_j)$. If a point (x_a, y_a) is below any such line segment, we flag it as 'not included', indicating that the upper concave envelope will not include this point. The overall scheme is complicated (see attached code) but runs significantly faster than the linear time method when batch sizes are greater than 32. An important downside, however is (a) the amount of *GPU* memory used for intermediate calculations and (b) the poor scaling (cubic) in terms of number of knots (which is βm , where β is the branching factor and m the number of knots) per EPF.

Dealing with different number of actions at each vertex and truncated points. Rather than removing points and padding them (to make each batch fit nicely in rectangular tensor), we maintain a 'mask' matrix which indicates that such a point is inactive. These points will not be used in computation of upper concave envelopes (both as potential points and as part of a line segment). Furthermore, this scheme makes it convenient to truncate points (simply mask those truncated points out and perform interpolation to get the new point on at $\tau(s')$).

G.4 Training Only Using Decreasing Portions of EPF

Kearning the increasing portion of an EPF is not useful, since points there are Pareto dominated. Hence, when extracting $\tilde{\pi}$ we will never select those points. For example in Figure 8b and 8c, if there were parts of the EPF in the yellow regions, they would not matter since the leader would select the maximum point with x-coordinate at 0 instead.

If we instead consider a slight variation of the EPF $U_s : \mathbb{R} \mapsto \mathbb{R} \cup \{-\infty\}$ that gives the maximum leader payoff given the follower gets a payoff of *at least* μ_2 (rather than exactly μ_2). This slight change ensures that EPFs are never increasing, while keeping all of the properties we proved earlier on. Omitting the increasing portions saves us from wasting any of the *m* knots on the increasing portions, and instead focus on the decreasing portion (where there is a real trade-off between payoffs between P₁ and P₂. One example of this is shown in Figure 4b and Figure 4c, where we showed the 'true' EPF and the modified EPFs that we use for our experiments.

We describe what happens concretely. Let \tilde{U}'_s be computed based on (1) with its representation given by the set of knots $\{(x_1, y_1), \ldots, (x_k, y_k)\}$, assumed to be sorted in ascending order of x-coordinates, and where $x_1 = \underline{V}(s)$. Let the $j = \arg \max_i y_i$. Then, the set of points which we use for training is the modified set $\{(\underline{V}(s), y_j), (x_j, y_j), (x_{j+1}, y_{j+1}), \ldots, (x_k, y_k)\}$.

G.5 Sampling of Training Trajectories

One of the design decisions in FVI is how one should sample states, or trajectories. In the single-player setting, it is commonplace to use some form of ϵ -greedy sampling. In our work, we use an even simple sampling scheme which takes actions uniformly at random.

There are a few exceptions. For RC sampled *states* uniformly at random. This was made possible because the game was small and we could enumerate all states. The implication is that each leaf is sampled much more frequently than from actual trajectories. Our experience is that since the game is small, getting samples from uniform trajectories should still work well. For TANTRUM, the game has a depth of 50. In many cases, states in the middle are not learning anything meaningful because their children EPFs have not been learned well. As such, we adopt a 'layered' approach, where initially we only allow for states s at most d_{max} to be added to the replay buffer, d_{max} is gradually increased as training goes on. This helps EPFs to be learned for states deeper in the tree first before their parents. We find that for TANTRUM (the non-featurized version with n = 25), this was essential to get stable learning of EPFs (recall that in our setting, TANTRUM has a size of roughly ~ 3^{25} and a depth of 50. A uniform trajectory leaves some states to be sampled with probability $1/2^{50}$). We start off at $d_{\text{max}} = 20$ and reduce d_{max} by 1 for every 50000 epochs. For other games, uniform trajectories work well enough since the game is not too deep.

H Additional Details on Experimental Setup

H.1 Environment Details

For all our experiments, the network is a multilayer fully connected network of width 128, depth 8, ReLU activations and number of knots m = 8. We used the PyTorch library (Paszke et al. 2017) and a GPU to accelerate training. No hyperparameter tuning was done.

H.2 RC Map Generation Details

Maps were generated with each reward map being drawn independently from a log Gaussian process (with query points given by the (x, y) coordinates on the grid). We use the square-exponential kernel, a length scale of 2.0 and a standard deviation of 0.1. This way of generating maps was to encourage spatial smoothness in rewards for more realism. Figure 10 give examples of maps generated using this procedure. From the figures, one cans see that good regions for P₁ may not be good for P₂ and vice versa.

H.3 TANTRUM Generation

For [TANTRUM], the values of q_2 were chosen (somewhat arbitrarily) to be in $\{1.5, 2.1, 3.4, 5.1, 6.7\}$.

H.4 Training Hyperparameters

We used the Adam optimizer (Kingma and Ba) with AMSGrad (Reddi et. al.) with a learning rate of 1e-5 (except for RC, where we found using a learning rate of 1e-4 was more suitable). We use the implementation provided in the PyTorch (Paszke et al. 2017) library. The replay buffer was of a size of 1M. The minibatch size was set to 128. The target network's parameters was updated once every 2000 training epochs.

H.5 Number of Epochs and Termination Criterion

Unfortunately, it is very rare to have a small loss for every single state. Hence, we terminate training at a fixed iteration. The number of epochs are given in Table 1.



Figure 10: Left to right: An example of reward maps used for P_1 and P_2 in RC+.

	Num Epochs	State sampling method	Prioritized replay	Traj. frequency
RC	2M	Random state	N/A	N/A
TANTRUM	4M	Unif. trajectory, layered	Yes	10
Feat. TANTRUM	2.7M	Uniform trajectory	Yes	10
RC+	1.7M	Uniform trajectory	Yes	20

Table 1: Differences in experimental setups over each game. Traj. frequency refers to how many epochs before we sample a new trajectory.

I Qualitative Discussion of Optimal Strategies in TANTRUM

We explore TANTRUM in the special case when $q_1 = 1$ and $q_2 > 1$. Intuitively, we should 'use' as many threats as possible. That is achieved by the leader committing to (-1, -1) for all future states. If P_1 does that from the beginning, it will give $P_2 - n$ (the number of times the stage game is repeated) payoff to each player. Naturally, one upper-bound on how much P_1 can get is n/q_2 , that is, P_2 chooses to accede on average of n/q_2 times per playthrough. P_1 cannot possible get more since that would lead to to P_2 losing more than n (which is the worst possible threat the leader can make from the beginning).

In our experiments, this was indeed true, and can be achieved by the following strategy. Let π be such that (a) the P₁ plays to (0,0) at all leader vertices and P₂ accedes for the first $j = \lfloor n/q_2 \rfloor$ stages with probability 1. At the j + 1-th vertex, it plays a mixed strategy (or rather, it receives the recommendation to mix) strategies, with probability $(n - jq_2)/q_2$ it accedes. Clearly, the expected payoff for P₂ is -n, and P₂ cannot do any better.

However, this result does not hold for all settings, typically when n is small. This is because we need to consider the threat is strong enough at every stage. Consider the case where n = 3 and $q_2 = 2$. Our derivation suggests that at the first stage, the follower accedes with probability 1 and at the second stage, it accedes with probability 0.5.

At the first stage, P_2 is indeed incentivized to accede. since if it will suffer from -3 if not, since acceding yields -2 payoff, which when combined with the expected payoff of -1 in the future, is equivalent to the threat of -3. At the second stage, it is just barely incentive compatible for the follower to accede. Specifically, if the player accedes, it will receive a payoff of -4 (it has already accumulated -2 from the previous stage). On the other hand, the grim trigger threat gives a payoff of only -4 (-2 from the past and -2 from the future). Hence, after receiving the recommendation P_2 is just incentivized to not deviate. However, when q_2 is increased by just a little (say to 2.1), this incentive is not sufficient. The **future** threat from not acceding is -2, but the follower already loses 2.1 from acceding. ⁸

In general, for our derived bound to be tight, we will require

$$\underbrace{n - \lfloor n/q \rfloor}_{\text{rest from origination}} \ge \underbrace{q}_{\text{rest from organizer this time rest}}$$

threat from grim trigger cost from acceding this time round

that is, at the last accede recommendation (possibly with some probability), we still have enough rounds remaining as threats to maintain incentive compatibility. Technically we require this for all previous rounds; however this condition being satisfied for the last round implies that it is satisfied for all previous rounds.

⁸This phenomena is very similar to the difference between a coarse correlated equilibrium and a regular CE. The difference is that a player has to decide to deviate before or after receiving its recommended actions.

We can also see from this discussion that in this repeated setting, it is always beneficial to recommend accede to the follower higher up the tree; this way, there is more room for the leader to threaten the follower with future (-1, -1) actions.

Featurized TANTRUM As far as we know, there is no simple closed form solution for featurized TANTRUM.