## A. Proof of Theorem 2

The key to the proof is to bound how much this naive normalization changes the point $x$. Let $(x^*, z^*)$ be the result of projecting $(x, z)$ into the optimal set $S$.

**Lemma.** *Let $x'$ be the result of normalizing $x$ according to the given scheme, and $i$ be an information set at depth $d$ (with the root defined to be at depth 0. Then we have $|x'_i - x^*_i| \le \varepsilon d\sqrt{n}$.*

*Proof.* By induction on the sequence-form strategy tree for player $x$, starting at the root. At the root node $i = 0$, the claim is clearly true because $x_0 = 1$ in any feasible solution $x$. Now consider any information set with parent $x_{i_0}$ and children $x_i := (x_{i_1}, \dots, x_{i_k})$ at depth $d$. From the theorem statement, we have $\|x_i - x^*_i\|_2 \le \varepsilon$, and since $x^*$ is feasible, we have $\sum_{j=1}^k x^*_{i_k} = x^*_{i_0}$. It follows that

$$\left|\sum_{j=1}^k x_{i_k} - x'_{i_0}\right| \le \sum_{j=1}^k \left|x_{i_k} - x^*_{i_k}\right| + \left|x'_{i_0} - x^*_{i_0}\right| \le \varepsilon\sqrt{k} + \varepsilon(d-1)\sqrt{n} \le \varepsilon d\sqrt{n}$$

by triangle inequality and inductive hypothesis, and noting that $k \le n$. But the normalization acts by picking $x'_i$ so that $\sum_{j=1}^k x'_{i_k} = x'_{i_0}$, and it moves all the $x_{i_k}$s in the same direction; thus, each one can move by at most $\varepsilon d\sqrt{n}$, completing the induction. □

With this lemma in hand, we now prove the theorem.

*Proof of Theorem.* Since $d \le n$ (each depth must have at least one information set), it follows from the lemma that $\|x' - x^*\|_2 \le \varepsilon n^2$. But the best response function $\min_y x^T Ay$ (with feasibility constraints on $y$) is a pointwise minimum of Lipschitz functions $x^T v$ for each $v = Ay$ and $y$ feasible, hence itself Lipschitz, with Lipschitz constant

$$\max_y \|Ay\|_2 \le \max_y \|Ay\|_1 \le \|A\|_1 \max_y \|y\|_\infty = \|A\|_1 \le n^2\|A\|_\infty.$$

where $\|A\|_1$ is the sum of the magnitudes of the nonzero entries of $A$. The desired theorem follows. □

## B. Another Example of the Utility of Sparse Factorization

**Example 3.** Let $A$ be the $n \times (n+1)$ matrix given by $A = \begin{bmatrix} I_n & 0 \end{bmatrix} + \begin{bmatrix} 0 & I_n \end{bmatrix}$, where $I_n$ is the $n \times n$ identity, and 0 is a column vector of zeros. So, $A$ is the matrix whose $(i, j)$ entry is 1 exactly when $j = i$ or $j = i + 1$. By direct computation, the SVD of this matrix is $A = U\Sigma V^T$ where $U$ and $V$ are *fully dense*, and the SVD is unique (in the usual sense, that is, up to signs and permutations) since all the singular values are. Thus, taking an SVD would have the result of *increasing* the number of nonzeros from $2n$ to $\Theta(n^2)$, which is the opposite of what we want. Thus, although in this case there will not be a good sparse factorization, using SVD make the problem worse.

## C. Benchmark Games in Experiment 1

We tested on the following benchmark games from the literature:

- *Leduc poker* (Southey et al., 2005) is a small variant of poker, played with one hole card and three community cards.
- *Battleship* (Farina et al., 2019) is the classic targeting game, with two parameters: $m$ is the number of moves (shots) a player may take, and $n$ is the number of ships on the board. All ships have length 2. A player scores a point only for sinking a full ship.

- *Sheriff* (Farina et al., 2019) is a simplified Sheriff of Nottingham game, modified to be zero-sum, played between a *smuggler* and a *sheriff*. The smuggler selects a *bribe amount* $b \in [0, B]$ and a number of illegal items $n \in [0, N]$ to try to smuggle past the sheriff. The sheriff then decides whether to inspect. If the sheriff does not inspect the cargo, then the smuggler scores $n - b$. If the sheriff inspects and finds no illegal items ($n = 0$), then the smuggler scores 3. If the sheriff inspects, and $n > 0$, then the smuggler scores $-2n$. The smuggler has far more sequences than the sheriff in this game.

- *No-limit hold-em (NLH) river endgames* are endgames encountered by the poker-playing agent Libratus (Brown & Sandholm, 2017), using the action abstraction used by Libratus. They both begin on the last betting round, when all five community cards are known. The normalization of $\|A\|_\infty = 1$ means that in these endgames, a Nash gap of 1 corresponds to 0.075 big blinds. Due to the explicit storage of the payoff matrix in this experiment, only extremely small no-limit endgames can be tested. In particular, *endgame A* here is the same as *endgame 7* in the next experiment (with a finer abstraction), and *endgame B* is the same as *endgame A* except with the starting pot size doubled to make the game smaller.