Learning Properties in Simulation-Based Games

Cyrus Cousins
UMass Amherst
Massachusetts, USA
cbcousins@umass.edu

Bhaskar Mishra
University of Florida
Florida, USA
bmishra1@ufl.edu

Enrique Areyan Viqueira
enriqueareyan.com
Washington, USA
enrique3@gmail.com

Amy Greenwald
Brown University
Rhode Island, USA
amy_greenwald@brown.edu

ABSTRACT

Empirical game-theoretic analysis (EGTA) is primarily concerned with learning equilibria of simulation-based games. Recent statistical approaches have tackled this problem by first learning a uniform approximation of the game's utilities, and then applying precision-recall theorems: i.e., all equilibria of the true game are approximate equilibria in the estimated game, and vice-versa. In this work, we generalize this approach to all game properties that are well-behaved (i.e., Lipschitz continuous in utilities), including regret (which defines Nash and correlated equilibria), adversarial values, power-mean welfare, and Gini social welfare. We show that, given a well-behaved welfare function, while optimal welfare is well-behaved, the welfare of optimal (i.e., welfare-maximizing or minimizing) equilibria is not well behaved. We thus define a related property based on a Lagrangian relaxation of the equilibrium constraints that is well behaved. We call this property Λ-stable welfare. As determining the welfare of an optimal equilibrium is an essential step in computing the price of anarchy, we conclude with a discussion of an alternative, more stable notion of anarchy based on Λ-stable welfare, which we call the anarchy gap.

KEYWORDS

Empirical Game-Theoretic Analysis; Price of Anarchy; Estimation

ACM Reference Format:
Cyrus Cousins, Bhaskar Mishra, Enrique Areyan Viqueira, and Amy Greenwald. 2023. Learning Properties in Simulation-Based Games. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

In recent years, empirical game-theoretic analysis (EGTA) has emerged as a powerful tool by which to analyze multiagent systems [2, 22, 23, 28], particularly when only a simulator of the game is available, rather than a precise description of the system, perhaps because of complicated stochastic dynamics. Such systems are called *simulation-based games*, or *black-box* games, and their empirical counterparts, which are derived from simulation data, are called *empirical games*.

Some properties of interest in traditional game-theoretic analysis include the set of equilibria,[1] the maximal welfare, the price of anarchy, etc. EGTA is concerned with characterizing the properties of simulation-based games, which is typically done by learning those properties in the corresponding empirical games. Statistical EGTA [2, 3, 22, 26, 30] is intended to give practitioners tools by which they can then test hypotheses like, the set of equilibria of an empirical game $\hat{\Gamma}$ coincides with the set of equilibria of the corresponding simulation-based game Γ.

Simulation-based games are noisy by their very nature. Indeed, multiple simulation queries are necessary before a practitioner can feel confident they have produced an accurate empirical game: i.e., accurate estimates of the players' utilities at all strategy profiles. Even when a game itself is well-estimated (e.g., in the sense of PAC-learning [24]), it is still not obvious how to derive a guarantee for a game property from one about the game's accuracy, as game properties are complex, non-linear functions of a game's utilities.

Areyan, *et al.* [2] and Tuyls, *et al.* [22] provide finite-sample guarantees on the estimated equilibria of simulation-based games, starting from the notion of a uniform approximation. A *uniform approximation* of a game is one in which all utilities are estimated to within the same error, simultaneously. Building on the work of Vorobeychik [26], who established guarantees in the case of infinitely many samples, the aforementioned authors prove that when an empirical game $\hat{\Gamma}$ uniformly approximates a (true) game Γ, all equilibria in Γ are approximate equilibria in $\hat{\Gamma}$, and all equilibria in $\hat{\Gamma}$ are approximate equilibria in Γ. They then develop various sampling algorithms based on standard concentration inequalities (e.g., Hoeffding [14], Bennett [7]) to learn uniform approximations of simulation-based games, and thereby estimate the equilibria of these games with finite sample guarantees.

In this paper, we seek methods to learn properties of simulation-based games beyond equilibria. Following the techniques developed in Areyan Viqueira *et al.* [2], we ask: what other game-theoretic properties of interest can be well-approximated given only a uniform approximation? We begin by studying a few focal game-theoretic properties: power-mean welfare, Gini-social welfare, adversarial value, and regret. We are specifically interested in the extremal values of these properties, e.g., best (optimal) and worst (pessimal) welfare values, and the outcomes that realize them, both of which we show are amenable to statistical EGTA methodology: i.e., can be estimated with finite-sample guarantees.

[1] All mention of equilibria in this paper refer to Nash equilibria.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

We then give several examples of properties—extreme equilibria (i.e., optimal and pessimal), the price of anarchy [15], and the price of stability [1, 21]—which are not amenable to statistical EGTA methodology. Indeed, naïve estimation of extreme equilibria via a uniform approximation can be arbitrarily incorrect. The prices of anarchy and stability are likewise inapproximable, because they are functions of extreme equilibria. In light of these observations, we introduce a relaxation of extreme equilibria, which we call Λ-*stable outcomes*. These outcomes are defined based on an extreme value of welfare discounted by the distance from an equilibrium, i.e., Λ acts like a Lagrangian. Unlike extreme equilibria, this new concept *is* amenable to statistical EGTA methodology.

To illustrate the utility of Λ-stable outcomes, we investigate the price of anarchy, i.e., the ratio between the best and worst welfare, the latter restricted to values at equilibria. We argue that even if the worst equilibrium were replaced by a pessimal Λ-stable outcome, we still could not estimate the price of anarchy, because, if the worst equilibrium welfare is near zero, the sensitivity to the best welfare outcome becomes infinite. We thus define an alternative concept, the *anarchy gap*, which instead measures the *difference* between the best welfare and the maximally dissonant Λ-stable outcome (a relaxation of a worst welfare equilibrium). This new concept is amenable to statistical EGTA methodology, as we show experimentally. An analogous story holds for the price of stability.

2 LEARNING EXTREMAL PROPERTIES

We begin by presenting our framework for approximating game-theoretic properties in simulation-based games. Doing so requires two pieces of technical machinery: the notion of uniform approximation and that of Lipschitz continuity. Given this machinery, it is immediate that if a game property is Lipschitz continuous in utilities, then it is *well behaved*, meaning it can be learned by any algorithm that produces a uniform approximation of the game.

In this work, we emphasize *extremal* game properties, such as the optimal or pessimal welfare (the latter being relevant, for example, when considering the price of anarchy [8]). We are interested in learning not only the *values* of these extremal properties, but further the strategy profiles that generate those values: i.e., the arguments that realize the solutions to the optimization problems that extremal properties define. We refer to these arguments as *witnesses* of the corresponding property. For example, an equilibrium is a witness of the regret property: i.e., it is a strategy profile that minimizes regret (and thereby attains an extremum).

Given an approximation of one game by another, there is not necessarily a connection between their extremal properties. For example, there may be equilibria in one game with no corresponding equilibria in the other, as small changes to the game can add or remove equilibria. Nonetheless, finding the equilibria of a uniform approximation of a game is sufficient for finding the approximate equilibria of the game itself [4]. The main result of this section is to generalize this result beyond regret, which defines equilibria, to all well-behaved game properties. We thus explain the aforementioned result as a consequence of a more general theory.

2.1 Standard Game Theory

We begin by defining games and several focal properties of games.

Definition 2.1 (Normal-Form Game). A *normal-form game* (NFG) $\Gamma \doteq \langle P, \{S_p\}_{p\in P}, \boldsymbol{u}\rangle$ consists of a set of players P, with *pure strategy set* S_p available to player $p \in P$. We define $\boldsymbol{S} \doteq S_1 \times \cdots \times S_{|P|}$ to be the *pure strategy profile space*, and then $\boldsymbol{u} : \boldsymbol{S} \to \mathbb{R}^{|P|}$ is a vector-valued utility function (equivalently, a vector of $|P|$ scalar utility functions $\boldsymbol{u}_p$).

Given an NFG Γ with finite S_p for all $p \in P$, we denote by $S_p^\diamond$ the set of distributions over S_p. This set is called player p's *mixed strategy set*. We define $\boldsymbol{S}^\diamond = S_1^\diamond \times \cdots \times S_{|P|}^\diamond$ to be the *mixed strategy profile space*, and then, overloading notation, we write $\boldsymbol{u}(\boldsymbol{s})$ to denote the expected utility of a mixed strategy profile $\boldsymbol{s} \in \boldsymbol{S}^\diamond$. We denote the *mixed game* that comprises mixed strategies $\boldsymbol{S}^\diamond$ by $\Gamma^\diamond$.

We call two NFGs with the same player sets and strategy profile spaces *compatible*. In this paper, we focus on uniform approximations of one NFG by another compatible one, by which we mean approximating one game's *utilities* by another's. Thus, we define the game properties of interest in terms of $\boldsymbol{u}$ rather than Γ, as the players and their strategy sets are usually clear from context. Likewise, we write $\boldsymbol{u}^\diamond$, rather than $\Gamma^\diamond$. We assume a NFG $\boldsymbol{u}$ in the definitions that follow.

Definition 2.2 (Property). A *property* of a game $\boldsymbol{u}$ is a functional mapping an index set $\mathcal{X}$ and utilities $\boldsymbol{u}$ to real values: i.e., $f : \mathcal{X} \times (\boldsymbol{S} \to \mathbb{R}^{|P|}) \to \mathbb{R}$.

In this work, two common choices for $\mathcal{X}$ are the set of pure and mixed strategy profiles $\boldsymbol{S}$ and $\boldsymbol{S}^\diamond$, respectively. Another plausible choice is the set of pure strategies for just one player p, namely S_p.

In this section, we focus on four properties, which we find to be well behaved: power-mean welfare, Gini social welfare, adversarial values, and regret.

Definition 2.3 (Power-Mean Welfare). Given strategy profile $\boldsymbol{s}$, power $\rho \in \mathbb{R}$, and stochastic weight vector $\boldsymbol{w} \in \triangle^{|P|}$, i.e., $\boldsymbol{w} \in \mathbb{R}_{0+}^{|P|}$ s.t. $\|\boldsymbol{w}\|_1 = 1$, the ρ-*power-mean welfare*[2] at $\boldsymbol{s}$ is defined as $\mathrm{W}_{\rho,\boldsymbol{w}}(\boldsymbol{s};\boldsymbol{u}) \doteq \sqrt[\rho]{\boldsymbol{w}\cdot\boldsymbol{u}(\boldsymbol{s})^\rho}$. When $\boldsymbol{w}$ is left unspecified (i.e., $\mathrm{W}_\rho(\boldsymbol{s};\boldsymbol{u})$), we use as default values $\boldsymbol{w} \doteq \left(\frac{1}{|P|}, \ldots, \frac{1}{|P|}\right)$.

A few special cases of ρ-power mean welfare are worth mentioning. When $\rho = 1$, power-mean welfare corresponds to utilitarian welfare, while when $\rho = -\infty$, power-mean welfare corresponds to egalitarian welfare. Finally, when $\rho = 0$, taking limits yields $\mathrm{W}_0(\boldsymbol{s};\boldsymbol{u}) \doteq \sqrt[|P|]{\prod_{p\in P}\boldsymbol{u}_p(\boldsymbol{s})}$, which defines Nash social welfare [16].

Definition 2.4 (Gini Social Welfare [29]). Given a strategy profile $\boldsymbol{s}$ and a decreasing stochastic weight vector $\boldsymbol{w}^\downarrow \in \triangle^{|P|}$, the *Gini social welfare* at $\boldsymbol{s}$ is defined as $\mathrm{W}_{\boldsymbol{w}^\downarrow}(\boldsymbol{s};\boldsymbol{u}) \doteq \boldsymbol{w}^\downarrow \cdot \boldsymbol{u}^\uparrow(\boldsymbol{s})$, where $\boldsymbol{u}^\uparrow(\boldsymbol{s})$ denotes the entries in $\boldsymbol{u}(\boldsymbol{s})$ in ascending sorted order.

In this definition, the weight vector $\boldsymbol{w}^\downarrow$ controls the trade-off between society's attitude towards well-off and impoverished players. We recover utilitarian welfare with $\boldsymbol{w}^\downarrow = (\frac{1}{|P|}, \frac{1}{|P|}, \ldots, \frac{1}{|P|})$, and egalitarian welfare with $\boldsymbol{w}^\downarrow = (1, 0, \ldots, 0)$.

The last two properties we consider, adversarial values and regret, relate to solution concepts. A player's *adversarial value* for playing a strategy is the value they obtain assuming worst-case behavior

[2] Power-mean welfare is more precisely defined as $\lim_{\rho'\to\rho} \sqrt[\rho']{\boldsymbol{w}\cdot\boldsymbol{u}(\boldsymbol{s})^{\rho'}}$, to handle the special cases when ρ is 0 or $\pm\infty$.

on the part of the other players: i.e., assuming all the other players were out to get them. A player's *regret* for playing one strategy measures how much they regret not playing another, fixing all the other players' strategies.

Definition 2.5 (Adversarial Values). A player p's *adversarial value* at strategy $\tilde{s} \in S_p^\diamond$ is defined as $\mathrm{A}_p(\tilde{s}; \boldsymbol{u}) \doteq \inf_{s \in S | s_p = \tilde{s}} \boldsymbol{u}_p(\boldsymbol{s})$.[3] A player p's *maximin value* is given by $\mathrm{MM}_p(\boldsymbol{u}) \doteq \sup_{\tilde{s} \in S_p^\diamond} \mathrm{A}_p(\tilde{s}; \boldsymbol{u}) = \sup_{\tilde{s} \in S_p^\diamond} \inf_{s \in S | s_p = \tilde{s}} \boldsymbol{u}_p(\boldsymbol{s})$. A strategy $\tilde{s}$ is *ε-maximin optimal* for player p if $\mathrm{A}_p(\tilde{s}; \boldsymbol{u}) \geq \mathrm{MM}_p(\boldsymbol{u}) - \varepsilon$.

Definition 2.6 (Regret). Fixing a player p and a strategy profile $\boldsymbol{s} \in \boldsymbol{S}$, we define $\mathrm{Adj}_{p,\boldsymbol{s}} \doteq \{\boldsymbol{t} \in \boldsymbol{S} \mid \boldsymbol{t}_q = \boldsymbol{s}_q, \forall q \neq p\}$: i.e., the set of adjacent strategy profiles, meaning those in which the strategies of all players $q \neq p$ are fixed at $\boldsymbol{s}_q$, while p's strategy may vary. Player p's *regret* at $\boldsymbol{s} \in \boldsymbol{S}$ is then defined as $\mathrm{Reg}_p(\boldsymbol{s}; \boldsymbol{u}) \doteq \sup_{\boldsymbol{s}' \in \mathrm{Adj}_{p,\boldsymbol{s}}} \boldsymbol{u}_p(\boldsymbol{s}') - \boldsymbol{u}_p(\boldsymbol{s})$, with $\mathrm{Reg}(\boldsymbol{s}; \boldsymbol{u}) \doteq \max_{p \in P} \mathrm{Reg}_p(\boldsymbol{s}; \boldsymbol{u})$.

Note that $\mathrm{Reg}_p(\boldsymbol{s}; \boldsymbol{u}) \geq 0$, since player p can deviate to any strategy, including $\boldsymbol{s}_p$ itself. Hence, $\mathrm{Reg}(\boldsymbol{s}; \boldsymbol{u}) \geq 0$. A strategy profile $\boldsymbol{s} \in \boldsymbol{S}$ that has regret at most $\varepsilon \geq 0$ is a called an *ε-Nash equilibrium* [17]: i.e., $\boldsymbol{s}$ is an ε-Nash equilibrium if and only if $0 \leq \mathrm{Reg}(\boldsymbol{s}; \boldsymbol{u}) \leq \varepsilon$.

2.2 Lipschitz Continuous Game Properties

Next, we define Lipschitz continuity, and show that the aforementioned game properties are all Lipschitz continuous in utilities.

Definition 2.7 (Lipschitz Property). Given $\lambda \geq 0$, a *λ-Lipschitz property* is one that is λ-Lipschitz continuous in utilities: i.e.,

$$\left\| f(\cdot; \boldsymbol{u}) - f(\cdot; \boldsymbol{u}') \right\|_\infty \doteq \sup_{x \in \mathcal{X}} \left| f(x; \boldsymbol{u}) - f(x; \boldsymbol{u}') \right| \leq \lambda \left\| \boldsymbol{u} - \boldsymbol{u}' \right\|_\infty \ ,$$

for all pairs of compatible games $\boldsymbol{u}$ and $\boldsymbol{u}'$.

To show that the game properties of interest are Lipschitz properties, we instantiate $\mathcal{X}$ and the property f in Definition 2.7 as follows:

(1) Power-Mean Welfare: Let $\mathcal{X} = \boldsymbol{S}$ and $f(\boldsymbol{s}; \boldsymbol{u}) = \mathrm{W}_{\rho, \boldsymbol{w}}(\boldsymbol{s}; \boldsymbol{u})$, for some $\rho \in \mathbb{R}$ and $\boldsymbol{w} \in \triangle^{|P|}$.
(2) Gini Social Welfare: Let $\mathcal{X} = \boldsymbol{S}$ and $f(\boldsymbol{s}; \boldsymbol{u}) = \mathrm{W}_{\boldsymbol{w}^\downarrow}(\boldsymbol{s}; \boldsymbol{u})$, for some $\boldsymbol{w}^\downarrow \in \triangle^{|P|}$.
(3) Adversarial Values: For $p \in P$, let $\mathcal{X} = S_p$ and $f_p = \mathrm{A}_p$: i.e., $f_p(\tilde{s}; \boldsymbol{u}) = \mathrm{A}_p(\tilde{s}; \boldsymbol{u})$.
(4) Regret: Let $\mathcal{X} = \boldsymbol{S}$ and $f = \mathrm{Reg}$: i.e., $f(\boldsymbol{s}; \boldsymbol{u}) = \mathrm{Reg}(\boldsymbol{s}; \boldsymbol{u})$.

Results on gradients and Lipschitz constants for power-mean welfare can be found in Beliakov et al. [6] and Cousins [10, 11, 12]. In short, whenever $\rho \geq 1$, power-mean welfare is Lipschitz continuous with $\lambda = 1$. It is also $\max_{p \in P} \boldsymbol{w}_p^{1/\rho}$-Lipschitz continuous for $\rho \in [-\infty, 0)$, but it is Lipschitz discontinuous for $\rho \in [0, 1)$.

The Lipschitz constants of the latter three properties can be derived using a "Lipschitz calculus," which is a straightforward consequence of the definition of Lipschitz continuity.

Theorem 2.8 (Lipschitz Calculus). *The following rules hold:*

(1) Linear Combination: If $g_{1:m}$ are $\lambda_{1:m}$-Lipschitz, all w.r.t. the same two norms, and $\boldsymbol{w} \in \mathbb{R}^m$, then the function $x \mapsto \boldsymbol{w} \cdot g(x)$ is $\sum_{i=1}^m \lambda_i |\boldsymbol{w}_i|$ -Lipschitz.
(2) Composition: If $h : A \to B$ and $g : B \to C$ are λ_h- and λ_g-Lispchitz w.r.t. norms $\|\cdot\|_A$, $\|\cdot\|_B$, & $\|\cdot\|_C$, then $(g \circ h) : A \to C$ is $\lambda_h \lambda_g$ Lispchitz w.r.t. $\|\cdot\|_A$ and $\|\cdot\|_C$.
(3) The infimum and supremum operations are 1-Lipschitz: i.e., if for all $x \in \mathcal{X}$, $f(x; \boldsymbol{u})$ is λ-Lipschitz in $\boldsymbol{u}$ then $\inf_{x \in \mathcal{X}} f(x; \boldsymbol{u})$ is also λ-Lipschitz in $\boldsymbol{u}$. Likewise, for the supremum.

By Theorem 2.8, Gini social welfare and adversarial value are both 1-Lipschitz, while regret is 2-Lipschitz. The interested reader is invited to consult the supplementary material for details.

The property f that computes a convex combination of utilities is 1-Lipschitz by the linear combination rule (Theorem 2.8), because utilities are 1-Lipschitz in themselves. Consequently, any findings about the Lipschitz continuity of game properties immediately apply to games with mixed strategies, as any λ-property g of a game $\boldsymbol{u}$ may be composed with f to arrive at a property $g \circ f$ of the mixed game $\boldsymbol{u}^\diamond$, which, by the composition rule (Theorem 2.8), remains λ-Lipschitz. In sum, all four of our focal properties are Lipschitz continuous in utilities.

Uniform Approximations of Game Properties. Next we observe that Lipschitz properties of a game are well behaved, and thus can be well-estimated by a uniform approximation of the game. Thus, all four of our focal game-theoretic properties are well behaved. A game $\boldsymbol{u}$ is an *ε-uniform approximation* of a compatible game $\boldsymbol{u}'$ if

$$\left\| \boldsymbol{u} - \boldsymbol{u}' \right\|_\infty \doteq \sup_{p \in P, \boldsymbol{s} \in \boldsymbol{S}} \left| \boldsymbol{u}_p(\boldsymbol{s}) - \boldsymbol{u}'_p(\boldsymbol{s}) \right| \leq \varepsilon \ .$$

In other words, in a uniform approximation of one NFG by another, the bound between utility deviations in the two games holds uniformly over *all* players and strategy profiles.

Our first observation, which follows immediately from the definitions of Lipschitz continuity and uniform approximation, characterizes well-behaved game properties:

Observation 2.1. *Let $\lambda, \varepsilon \geq 0$. If f is λ-Lipschitz and $\|\boldsymbol{u} - \boldsymbol{u}'\|_\infty \leq \varepsilon$, then $\|f(\cdot; \boldsymbol{u}) - f(\cdot; \boldsymbol{u}')\|_\infty \leq \lambda\varepsilon$. Furthermore, if $\mathcal{X} = \{x\}$ is a singleton, then $|f(x; \boldsymbol{u}) - f(x; \boldsymbol{u}')| \leq \lambda\varepsilon$.*

Equivalently, if f is λ-Lipschitz and $\|\boldsymbol{u} - \boldsymbol{u}'\|_\infty \leq \varepsilon/\lambda$, then $\|f(\cdot; \boldsymbol{u}) - f(\cdot; \boldsymbol{u}')\|_\infty \leq \lambda(\varepsilon/\lambda) = \varepsilon$. For example, regret is 2-Lipschitz, and thus can be ε-approximated, given an $\varepsilon/2$-uniform approximation. Likewise for adversarial values and variants of welfare, with their corresponding Lipschitz constants.

We visualize the above observation in the context of power-mean welfare in Figure 1. This plot was created using GAMUT [18] to generate a congestion game $\boldsymbol{u}$ with 3 players, 3 facilities, and utility range $[-1, 1]$. We uniformly drew 100 normal-form games from the collection of all 0.02-uniform approximations of this congestion game. For values of $\rho \in (-\infty, \infty) \cup \{\pm\infty\}$, we plotted the average across all 100 games of the supremum error across strategy profiles $\boldsymbol{s} \in \boldsymbol{S}$ between the estimate of the ρ-power-mean welfare at $\boldsymbol{s}$ via the given uniform approximation $\boldsymbol{u}'$ of $\boldsymbol{u}$ and the true ρ-power-mean welfare at $\boldsymbol{s}$ (i.e., $\sup_{\boldsymbol{s} \in \boldsymbol{S}} \left| \mathrm{W}_\rho(\boldsymbol{s}; \boldsymbol{u}) - \mathrm{W}_\rho(\boldsymbol{s}; \boldsymbol{u}') \right|$), using weights $\boldsymbol{w}_p = 1/|P|$ for all $p \in P$. We also plot the upper bounds on this supremum error implied by Observation 2.1, using the trivial upper bound of 2 (since the utility range is $[-1, 1]$) in cases where ρ renders W_ρ not Lipschitz continuous. Figure 1 shows that ρ-power-mean welfare is ε-well-estimated for $\rho > 1$, as expected, and is also

[3] Equipping players other than p with mixed strategies affords them no added power.

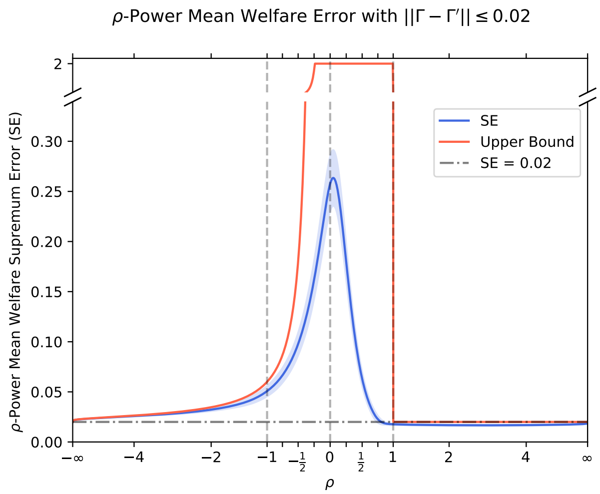

Figure 1: Average supremum error between the true ρ-power-mean welfare at each strategy profile of a congestion game and the corresponding estimates via 100 random 0.02-uniform approximations of the true game, with standard deviations shaded in blue, for $\rho \in (-\infty, \infty)$. A sigmoid transformation is applied to the x-axis yielding a domain of $(-\infty, \infty)$.

fairly well-estimated for $\rho < -1$. Since ρ-power-mean welfare is $\max_{p \in P} \boldsymbol{w}_p^{1/\rho}$-Lipschitz continuous for $\rho \in [-\infty, 0)$, we expect high error as $\rho \to 0^-$, and as it is Lipschitz discontinuous for $\rho \in [0, 1)$, we also expect high error in this region. Figure 1 is also consistent with these expectations.

Although power means are not Lipschitz-continuous for $\rho \in [0, 1)$, and may consequently be more difficult to estimate than Lipschitz-continuous properties, they are not necessarily inapproximable via uniform approximations. In particular, for $\rho = 0$, the difficulty with estimating the power mean stems from the fact that its derivative can grow in an unbounded fashion if even one of the game's utilities approaches zero. As we move away from this singularity, however, power-mean estimates do converge as the accuracy of uniform approximations increases, but the rate of convergence may be slow. In contrast, other interesting properties of games, like the price of anarchy and maximal welfare at an equilibrium, are not even continuous (let alone Lipschitz continuous), and thus no convergent estimator can exist for them.

2.3 Approximating Extremal Properties and their Witnesses

Using Observation 2.1, we now show[4] that the extrema of Lipschitz continuous properties (e.g., optimal and pessimal welfare) are also well behaved, and that this result extends to their witnesses, all of which remain approximately optimal in their approximate game counterparts. Specifically, we derive a two-sided approximation bound on the values of these extremal properties, and a dual containment result characterizing their witnesses. This result can be understood as a form of recall and (approximate) precision: recall, because the set of witnesses of the approximate game $\hat{\boldsymbol{u}}$ contains all the true positives (i.e., all the witnesses in the game $\boldsymbol{u}$); precision, because all false positives (witnesses in the approximate game $\hat{\boldsymbol{u}}$ that are not witnesses in the game $\boldsymbol{u}$) are nonetheless approximate witnesses in the game $\boldsymbol{u}$. Taking the property of interest to be regret, so that its minima correspond to Nash equilibria, this final result can be understood as a generalization of the precision and recall result obtained in Areyan Viqueira et al. [4].

[4]All omitted proofs can be found in the supplementary material.

Theorem 2.9 (Approximating Extremal Properties of Normal-Form Games). *Let $\varepsilon, \lambda, \alpha > 0$. Given a λ-Lipschitz property f, the following hold:*

(1) Approximately-optimal values*: $\left|\sup_{x \in \mathcal{X}} f(x; \boldsymbol{u}) - \sup_{x \in \mathcal{X}} f(x; \boldsymbol{u}')\right| \leq \lambda\varepsilon$ and $\left|\inf_{x \in \mathcal{X}} f(x; \boldsymbol{u}) - \inf_{x \in \mathcal{X}} f(x; \boldsymbol{u}')\right| \leq \lambda\varepsilon$.*

(2) Approximately-optimal witnesses*: if some $\hat{x} \in \mathcal{X}$ is α-optimal according to $\boldsymbol{u}'$, then $\hat{x}$ is $2\lambda\varepsilon + \alpha$-optimal according to $\boldsymbol{u}$: i.e., if $f(\hat{x}; \boldsymbol{u}') \geq \sup_{x \in \mathcal{X}} f(x; \boldsymbol{u}') - \alpha$, then*

$$f(\hat{x}; \boldsymbol{u}) \geq \sup_{x \in \mathcal{X}} f(x; \boldsymbol{u}) - 2\lambda\varepsilon - \alpha \ ,$$

and if $f(\hat{x}; \boldsymbol{u}') \leq \inf_{x \in \mathcal{X}} f(x; \boldsymbol{u}') + \alpha$, then

$$f(\hat{x}; \boldsymbol{u}) \leq \inf_{x \in \mathcal{X}} f(x; \boldsymbol{u}) + 2\lambda\varepsilon + \alpha \ .$$

Applying this theorem to welfare and adversarial values yields the following corollary.

Corollary 2.10. *Both of the following hold:*

(1) Welfare: Let W denote a λ-Lipschitz welfare function. Then

$$\left|\sup_{x \in \mathcal{X}} \mathrm{W}(x; \boldsymbol{u}) - \sup_{x \in \mathcal{X}} \mathrm{W}(x; \boldsymbol{u}')\right| \leq \lambda\varepsilon$$
$$\text{and} \left|\inf_{x \in \mathcal{X}} \mathrm{W}(x; \boldsymbol{u}) - \inf_{x \in \mathcal{X}} \mathrm{W}(x; \boldsymbol{u}')\right| \leq \lambda\varepsilon \ .$$

(2) Maximin-Optimal Strategies: If strategy $\tilde{s} \in S_p$ is α-maximin optimal for player p in $\boldsymbol{u}'$, so that $\mathrm{A}_p(\tilde{s}; \boldsymbol{u}') \geq \mathrm{MM}_p(\boldsymbol{u}') - \alpha$, then it is $2\varepsilon - \alpha$ maximin-optimal in $\boldsymbol{u}$, meaning $\mathrm{A}_p(\tilde{s}; \boldsymbol{u}) \geq \mathrm{MM}_p(\boldsymbol{u}) - 2\varepsilon - \alpha$.

If the extreme value is known, we obtain a stronger result:

Theorem 2.11 (Approximating Witnesses of Normal-Form Games). *If v^* denotes the target value of a λ-Lipschitz property f and $F_\alpha(\boldsymbol{u}) = \{x \mid |f(x; \boldsymbol{u}) - v^*| \leq \alpha\}$, then $F_0(\boldsymbol{u}) \subseteq F_{\lambda\varepsilon}(\boldsymbol{u}') \subseteq F_{2\lambda\varepsilon}(\boldsymbol{u})$.*

By applying this theorem to regret, for which the extreme value is known (i.e., at equilibrium, regret is 0), we recover the dual containment theorem of Areyan Viqueira et al. [4].

Corollary 2.12 (Approximating Equilibria in NFGs). *If $\mathrm{E}_\alpha(\boldsymbol{u}) \doteq \{x \in \boldsymbol{S} \mid \mathrm{Reg}(x; \boldsymbol{u}) \leq \alpha\}$ and $\mathrm{E}^\diamond_\alpha(\boldsymbol{u}) \doteq \{x \in \boldsymbol{S}^\diamond \mid \mathrm{Reg}(x; \boldsymbol{u}) \leq \alpha\}$, then $\mathrm{E}_0(\boldsymbol{u}) \subseteq \mathrm{E}_{2\varepsilon}(\boldsymbol{u}') \subseteq \mathrm{E}_{4\varepsilon}(\boldsymbol{u})$ and $\mathrm{E}^\diamond_0(\boldsymbol{u}) \subseteq \mathrm{E}^\diamond_{2\varepsilon}(\boldsymbol{u}') \subseteq \mathrm{E}^\diamond_{4\varepsilon}(\boldsymbol{u})$.*

3 LEARNING EXTREME EQUILIBRIA

In this section, we investigate extreme equilibria—the best (optimal) or the worst (pessimal)—measured in terms of the welfare achieved. We call the welfare-maximal equilibria maximally *consonant*, and the welfare-minimizing ones, maximally *dissonant*.

As it turns out, even for a Lipschitz continuous welfare function, neither the maximally consonant nor maximally dissonant equilibria, nor their values, is amenable to statistical EGTA methodology, because these properties are not Lipschitz continuous in utilities. Consequently, we define relaxations of these properties, that we call

	A	B
A	$\boldsymbol{\gamma, -\gamma}$	$-\gamma, \gamma$
B	$\gamma - c, -\gamma$	$\boldsymbol{c - \gamma, \gamma}$

Figure 2: Inapproximable extreme equilibria.

maximally consonant and maximally dissonant Λ-stable outcomes, which are similar in spirit to their equilibrium counterparts, but which are well behaved. Specifically, we replace the equilibrium constraints with their Lagrangian relaxations based on a tunable parameter $\Lambda \geq 0$. This change in the definitions does not alter the spirit of the properties, only the letter, because as Λ goes to infinity, the penalty for violating the constraints becomes infinite, so they are not violated. On the other hand, when Λ is finite, these definitions can be understood as permitting small oscillations around extreme equilibria, thus expanding the scope of acceptable play.

For extreme Λ-stable outcomes, we obtain a positive result, namely, for all games and for all $\varepsilon > 0$, there exists a sample size m s.t. we can ε-estimate them. However, having overcome what might have appeared to be the shortcoming that was preventing the estimation of extreme equilibria, we find ourselves facing a more serious stumbling block. In particular, although we can now derive a more satisfying bound on these extreme consonance and dissonance properties, given an ε-uniform approximation of a game, this bound grows with Λ. In other words, fixing the number of samples and letting Λ increase so as to produce solutions closer and closer to an extreme equilibrium yields a larger and larger confidence interval around the property's estimate. Alternatively, fixing ε, as Λ grows, ε-accurate estimation of the property requires more and more samples. This result is not entirely surprising, as we have effectively interpolated between two extreme cases: 0-stable outcomes, which are easy to estimate, and ∞-stable outcomes, which cannot be estimated.

Observation 3.1 (Extreme Utilitarian Welfare Equilibria are Inapproximable). *Consider the game family $\Gamma(\gamma)$ parameterized by $\gamma \in (0, 1)$, for any $c \geq 0$, depicted in Figure 2. For $\Gamma(-\gamma)$ and $\Gamma(\gamma)$ with corresponding utilities $\boldsymbol{u}_{-\gamma}, \boldsymbol{u}_{\gamma}$, it holds that, for all $\varepsilon \geq \frac{|\gamma|}{2}$, it holds that $\left\|\boldsymbol{u}_{-\gamma} - \boldsymbol{u}_{\gamma}\right\|_\infty \leq 2\gamma \leq \varepsilon$, but*

$$\left|\sup_{\boldsymbol{s}\in \mathrm{E}(\boldsymbol{u})} \mathrm{W}_1(\boldsymbol{s};\boldsymbol{u}) - \sup_{\boldsymbol{s}\in \mathrm{E}(\boldsymbol{u}_\gamma)} \mathrm{W}_1(\boldsymbol{u}_\gamma, \boldsymbol{s})\right| = c$$

$$\text{and} \left|\inf_{\boldsymbol{s}\in \mathrm{E}(\boldsymbol{u})} \mathrm{W}_1(\boldsymbol{s};\boldsymbol{u}) - \inf_{\boldsymbol{s}\in \mathrm{E}(\boldsymbol{u}_\gamma)} \mathrm{W}_1(\boldsymbol{u}_\gamma, \boldsymbol{s})\right| = c \ .$$

Proof. By construction, $\Gamma(-\gamma)$ is a $2\gamma \leq \varepsilon$-uniform approximation of $\Gamma(\gamma)$. When $\gamma < 0$, the column player plays A, to which the row player responds with A. The only, and thus both the best and worst equilibrium of Γ is (A, A), with utilitarian welfare 0. For $\gamma > 0$, by similar reasoning the only equilibrium is (B, B), which has utilitarian welfare c. The utilitarian welfare of the best (and worst) welfare equilibria in $\Gamma(-\gamma)$ and $\Gamma(\gamma)$, therefore, differ by $c > 0$: i.e., arbitrarily. □

Corollary 3.1. *For all $\varepsilon > 0$, there does not exist a finite sample size m that is sufficient to ε-estimate maximally consonant and maximally dissonant w.r.t. utilitarian welfare) equilibria, or their values.*

Next, we derive natural bounds on extreme equilibria, both from above and below, albeit loosely. We then present an unnatural, yet sufficient, condition under which the upper (lower) bound on the maximally dissonant (consonant) equilibrium value can be tightened. We write $\mathrm{MD}(\boldsymbol{u})$ and $\mathrm{MC}(\boldsymbol{u})$ to denote the set of maximally dissonant and maximally consonant equilibria of the game defined by $\boldsymbol{u}$. As all elements of these sets have the same values, we overload this notation, using it to denote both the set (i.e., the witnesses) and the value of all its elements. Our intended meaning should be clear from context.

Given a game $\boldsymbol{u}$, the maximally dissonant equilibrium value $\mathrm{MD}(\boldsymbol{u})$ is upper bounded by the maximally consonant one $\mathrm{MC}(\boldsymbol{u})$. To obtain a lower bound on $\mathrm{MD}(\boldsymbol{u})$, we consider the value of the maximally dissonant 2ε-equilibria in an ε-approximation $\boldsymbol{u}'$ of $\boldsymbol{u}$, which underestimates the maximally dissonant equilibrium value of $\boldsymbol{u}$. Similarly, to obtain an upper bound on $\mathrm{MC}(\boldsymbol{u})$, we consider the value of the maximally consonant 2ε-equilibria in $\boldsymbol{u}'$, which overestimates the maximally consonant equilibrium value of $\boldsymbol{u}$. Note that both these bounds can be arbitrarily loose, as the extreme 2ε-equilibria in $\boldsymbol{u}'$ may not correspond to extreme equilibria in $\boldsymbol{u}$ (see Observation 3.1). That said, we can tighten this bound if we assume that the welfare of *every* α-Nash equilibrium in $\boldsymbol{u}$ approximates (by $\alpha\gamma$ for some $\gamma > 0$) that of an exact Nash equilibrium (also in $\boldsymbol{u}$), as this (unreasonable) assumption enables accurate estimation of the welfare of extreme equilibria.

We now formalize these arguments. As usual, we fix two compatible games characterized by utility functions $\boldsymbol{u}$ and $\boldsymbol{u}'$, and we assume that $\boldsymbol{u}'$ is an ε-uniform approximation of $\boldsymbol{u}$.

Theorem 3.2 (Approximating Extreme Equilibria). *Fix $\varepsilon > 0$. Assume $\mathrm{W} : \boldsymbol{S} \times (\boldsymbol{S} \to \mathbb{R}^{|P|}) \to \mathbb{R}$ is monotonic non-decreasing so that $\boldsymbol{u} \leq \boldsymbol{v}$ component-wise implies $\mathrm{W}(\boldsymbol{s};\boldsymbol{u}) \leq \mathrm{W}(\boldsymbol{s};\boldsymbol{v})$, for all $\boldsymbol{s} \in \boldsymbol{S}$ and utilities $\boldsymbol{u}$ and $\boldsymbol{v}$. It holds that*

$$\inf_{\boldsymbol{s}\in \mathrm{E}_{2\varepsilon}(\boldsymbol{u}')} \mathrm{W}(\boldsymbol{s};\boldsymbol{u}' - \varepsilon) \leq \mathrm{MD}(\boldsymbol{u}) \leq \mathrm{MC}(\boldsymbol{u}) \leq \sup_{\boldsymbol{s}\in \mathrm{E}_{2\varepsilon}(\boldsymbol{u}')} \mathrm{W}(\boldsymbol{s};\boldsymbol{u}' + \varepsilon) \ .$$

Corollary 3.3. *In addition to the conditions of Thm. 3.2, suppose that there exists $\gamma > 0$ such that*

$$\forall \alpha > 0, \forall \boldsymbol{s} \in \mathrm{E}_\alpha(\boldsymbol{u}), \exists \boldsymbol{s}' \in \mathrm{E}(\boldsymbol{u}) \text{ s.t. } \left|\mathrm{W}(\boldsymbol{s};\boldsymbol{u}) - \mathrm{W}(\boldsymbol{s};\boldsymbol{u}')\right| \leq \alpha\gamma \ .$$

We may then refine the upper bound on $\mathrm{MD}(\boldsymbol{u})$ as

$$\mathrm{MD}(\boldsymbol{u}) \leq \inf_{\boldsymbol{s}\in \mathrm{E}(\boldsymbol{u}')} \mathrm{W}(\boldsymbol{s};\boldsymbol{u}' + \varepsilon) + 2\gamma\varepsilon \ ,$$

and similarly, the lower bound on $\mathrm{MC}(\boldsymbol{u})$ as

$$\sup_{\boldsymbol{s}\in \mathrm{E}(\boldsymbol{u}')} \mathrm{W}(\boldsymbol{s};\boldsymbol{u}' - \varepsilon) - 2\gamma\varepsilon \leq \mathrm{MC}(\boldsymbol{u}) \ .$$

The condition that $|\mathrm{W}(\boldsymbol{s};\boldsymbol{u}) - \mathrm{W}(\boldsymbol{s};\boldsymbol{u}')| \leq \alpha\gamma$ is essentially a variant of the ε-Δ *approximate stability* property introduced by Awasthi *et al.* [5], who find that games with certain structures, such as zero-sum games, satisfy these conditions. Interestingly, while they use the condition to bound the error induced by *computational* approximation, our usage demonstrates that it can also be used to control the error due to *statistical* approximation.

As the condition of Corollary 3.3 is difficult to justify, we propose an alternative quantity as a substitute measure for extreme equilibria, namely extreme Λ-stable outcomes, in which we replace the

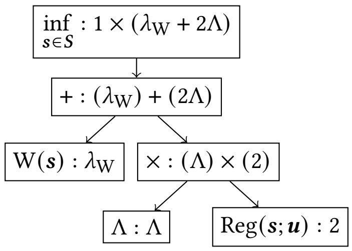

Figure 3: Abstract syntax tree depicting the derivation of the Lipschitz constant for $\mathrm{MD}_\Lambda(\boldsymbol{u})$.

equilibrium constraints with their Lagrangian relaxation. We now write $\mathrm{MD}_\Lambda(\boldsymbol{u})$ (resp. $\mathrm{MC}_\Lambda(\boldsymbol{u})$) to denote the maximally dissonant (resp. consonant) Λ-stable *outcomes* of the game defined by $\boldsymbol{u}$.

Definition 3.4. Given a game Γ with corresponding utilities $\boldsymbol{u}$, and $\Lambda \geq 0$, a *maximally dissonant* Λ*-stable outcome* is an element of

$$\mathrm{MD}_\Lambda(\boldsymbol{u}) \doteq \inf_{s \in S} \mathrm{W}(\boldsymbol{s}; \boldsymbol{u}) + \Lambda \, \mathrm{Reg}(\boldsymbol{s}; \boldsymbol{u}) \ .$$

Similarly, a *maximally consonant* Λ*-stable outcome* is an element of

$$\mathrm{MC}_\Lambda(\boldsymbol{u}) \doteq \sup_{s \in S} \mathrm{W}(\boldsymbol{s}; \boldsymbol{u}) - \Lambda \, \mathrm{Reg}(\boldsymbol{s}; \boldsymbol{u}) \ .$$

These new properties are Lipschitz continuous in utilities, as we argue next, and thus are better behaved than extreme equilibria, in that they can be well approximated via a uniform approximation. We also provide a proof that $\mathrm{MD}^*_\Lambda(\boldsymbol{u})$ and $\mathrm{MC}^*_\Lambda(\boldsymbol{u})$ are Lipschitz properties. These quantities are analogous to their counterparts, but are defined in terms of the *excess regret*

$$\mathrm{Reg}^*(\boldsymbol{s}; \boldsymbol{u}) \doteq \mathrm{Reg}(\boldsymbol{s}; \boldsymbol{u}) - \inf_{s' \in S} \mathrm{Reg}(\boldsymbol{s}'; \boldsymbol{u}) \ ,$$

which is relevant when the smallest regret is not zero. The Lipschitz constant of $\mathrm{Reg}^*(\boldsymbol{s}; \boldsymbol{u})$ is 4, as can be shown via similar techniques.

Theorem 3.5 (Approximating Extreme Λ-Stable Outcomes). *Assume* $\mathrm{W} : \boldsymbol{S} \times (\boldsymbol{S} \to \mathbb{R}^{|P|}) \to \mathbb{R}$ *is a* λ_{W}*-Lipschitz property. It holds that* $\mathrm{MD}_\Lambda(\boldsymbol{u})$ *is* $\lambda_{\mathrm{W}} + 2\Lambda$*-Lipschitz, which immediately implies that* $|\mathrm{MD}_\Lambda(\boldsymbol{u}) - \mathrm{MD}_\Lambda(\boldsymbol{u}')| \leq (\lambda_{\mathrm{W}} + 2\Lambda)\varepsilon$*. Likewise, for* $\mathrm{MC}_\Lambda(\boldsymbol{u})$*. It also holds that* $\mathrm{MD}^*_\Lambda(\boldsymbol{u})$ *is* $\lambda_{\mathrm{W}} + 4\Lambda$*-Lipschitz, which immediately implies that* $\left|\mathrm{MD}^*_\Lambda(\boldsymbol{u}) - \mathrm{MD}_\Lambda(\boldsymbol{u}')\right| \leq (\lambda_{\mathrm{W}} + 4\Lambda)\varepsilon$*. Likewise, for* $\mathrm{MC}^*_\Lambda(\boldsymbol{u})$.

Proof. We prove this theorem using the Lipschitz calculus presented in Theorem 2.8. Recall that regret is 2-Lipschitz. Moreover inf provides a 1-Lipschitz multiplicative factor. Finally, addition requires that we add the Lipschitz constants of the corresponding addends. We begin at the leaves in Figure 3, computing Lipschitz constants, and back those values up the syntax tree to arrive at $2(\lambda_{\mathrm{W}} + \Lambda)$ Lipschitz constant for $\mathrm{MD}_\Lambda(\boldsymbol{u})$. The inequality then follows via Observation 2.1.

The proofs in the remaining three cases — $\mathrm{MC}_\Lambda(\boldsymbol{u})$, $\mathrm{MD}^*_\Lambda(\boldsymbol{u})$, and $\mathrm{MC}^*_\Lambda(\boldsymbol{u})$ — are analogous. □

Discussion. Not only are extreme Λ-stable outcomes mathematically preferable to extreme equilibria because of their estimation properties, we believe they are also as, if not more, justifiable as descriptors or predictors of the play of a game. The standard motivation for considering welfare *at equilibrium* is that we expect ideal players to converge to equilibrium play. On the other hand, it is reasonable to expect non-ideal players to play near, but not exactly at, equilibria. By expanding the scope of play to include approximate equilibria, we are able to accommodate more realistic behavior. At the same time, the parameter Λ allows us to control the extent to which non-equilibrium play is acceptable. Taking $\Lambda = 0$ ignores equilibrium behavior entirely, while at the other extreme, letting $\Lambda \to \infty$, $\mathrm{MD}_\Lambda(\boldsymbol{u}) \to \mathrm{MD}(\boldsymbol{u})$ (likewise, $\mathrm{MC}_\Lambda(\boldsymbol{u}) \to \mathrm{MC}(\boldsymbol{u})$), because as Λ goes to infinity, the penalty for violating the constraints becomes infinite, so they are not violated.

We can interpret Λ-stable outcomes that deviate from equilibrium play as ranging from fully cooperative to fully non-cooperative. The welfare term is a cooperative one, where the players are colluding to make the world as wonderful (or terrible) a place as possible. The regret term, in contrast, is non-cooperative, but Λ permits some flexibility in behavior. In the case of a maximally dissonant (colluding) outcome, $\Lambda = 0$ implies the players are all playing pessimally: i.e., they are colluding to cause as much suffering to all players (including themselves) as possible. The opposite is true of maximally consonant (colluding) outcome, where players are colluding to create a utopia. On the other hand, as $\Lambda \to \infty$, the players move away from cooperative behavior towards rational equilibrium (i.e., non-cooperative) behavior. Indeed, Λ-stability allows us to model a range of cooperative and non-cooperative behaviors.

4 ESTIMATING THE PRICE OF ANARCHY

We now set out to investigate approximations of the price of anarchy [15] that follow from uniform approximations of games. The price of anarchy is a way of comparing the best-case welfare to the worst-case welfare at equilibrium. It is a measure of how bad things can get when there is no centralized control over agents' behavior, and the alternative assumption is worst-case equilibrium play.

Our first result is negative: even given a uniform approximation of a game, we cannot derive a satisfactory bound on its price of anarchy, as we cannot derive a satisfactory bound on the value of its maximally dissonant equilibria (see Obs. 3.1). Still, following Thm. 3.2, we can bound the price of anarchy both from above and below, albeit loosely from below, and then, under the same conditions proposed in that theorem, the lower bound can be tightened. But, as already noted, these conditions are difficult to justify, which leads us to propose an alternative quantity as a substitute measure for the traditional price of anarchy.

Our new measure of the price of anarchy alters the traditional measure in two ways. First, we replace the value of the maximally dissonant equilibrium with that of the maximally dissonant Λ-stable outcome, for some $\Lambda \geq 0$. Second, we replace division with subtraction. We refer to the traditional price of anarchy measure as the *anarchy ratio*, by analogy with competitive ratio, and our new measure as the *anarchy gap*. Although we use the term "gap" rather than "price," we note that price of anarchy is traditionally defined as a ratio, dividing one quantity by another; as such, it is dimensionless, and thus not technically a price. When you subtract the two quantities, however, the units are preserved, so that the resulting quantity remains measured in, say, *utils*, meaning units of utility, which interpreted in monetary terms would indeed be a price.

As in the case of extreme equilibria, the first change in the definition does not alter the spirit of the price of anarchy, only the letter. The second change, in contrast, is more material; it can, however, be justified mathematically. Replacing the division with a subtraction eliminates the possibility of dividing by a very small quantity. Together with the first change, this second change renders the anarchy gap Lipschitz continuous in utilities (Thm. 4.5). In contrast, the anarchy ratio is not Lipschitz continuous in utilities, even when it is defined in terms of Λ-stable outcomes. As a result, the anarchy gap is well behaved, and thus can be well-estimated using statistical EGTA methodology, while the anarchy ratio is not, and thus cannot.

Additionally, the price of anarchy was originally defined for complete-information games, so it was natural to extend the definition to incomplete-information games by computing the price of anarchy of the corresponding normal-form game, as is done in the study of the price of anarchy in auctions [20]. Computing the price of anarchy in this way, however, means dividing the 'expected value of the optimal social welfare' by 'the expected value of the maximally dissonant equilibrium value', rather than defining the price of anarchy as the expected value of 'the optimal social welfare divided by the maximally dissonant equilibrium value': i.e., taking expectations before dividing rather than after. Because of the linearity of expectations, the disparity between these two potential definitions the price of anarchy in games of incomplete-information disappears in a measure that is defined via subtraction rather than division: i.e., in the anarchy gap. We start our technical discussion by defining the price of anarchy [15].

Definition 4.1 (Anarchy Ratio). Given game Γ with utility function $\boldsymbol{u}$, the *anarchy ratio* is defined as

$$\mathrm{AR}(\boldsymbol{u}) \doteq \frac{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{u},\boldsymbol{s})}{\inf_{\boldsymbol{s}\in\mathrm{E}(\boldsymbol{u})} \mathrm{W}(\boldsymbol{u},\boldsymbol{s})} = \frac{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{u},\boldsymbol{s})}{\mathrm{MD}(\boldsymbol{u})} \enspace .$$

Likewise, the stability ratio[5] is defined as

$$\mathrm{SR}(\boldsymbol{u}) \doteq \frac{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{u},\boldsymbol{s})}{\sup_{\boldsymbol{s}\in\mathrm{E}(\boldsymbol{u})} \mathrm{W}(\boldsymbol{u},\boldsymbol{s})} = \frac{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{u},\boldsymbol{s})}{\mathrm{MC}(\boldsymbol{u})} \enspace .$$

As argued above, a uniform approximation $\boldsymbol{u}'$ of $\boldsymbol{u}$ implies a uniform approximation of the extreme welfare of $\boldsymbol{u}$ — the term in the numerator of the anarchy ratio. It does not, however, imply a uniform approximation of the welfare of an extreme equilibrium in $\boldsymbol{u}$ (see Observation 3.1) — the term in the denominator. Just as we bounded the values of extreme equilibria from above and below in Thm. 3.2, the next theorem bounds the anarchy ratio. These arguments are similar because the anarchy ratio is defined in terms of the maximally dissonant equilibrium value.

Theorem 4.2 (Approximating the Prices of Anarchy and Stability). *Assume as in Thm. 3.2. Then, so long as the numerator and denominator of each bound is positive,*

$$\max\left\{1, \frac{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{u}'-\varepsilon,\boldsymbol{s})}{\sup_{\boldsymbol{s}\in\mathrm{E}_{2\varepsilon}(\boldsymbol{u}')} \mathrm{W}(\boldsymbol{u}'+\varepsilon,\boldsymbol{s})}\right\} \le \mathrm{SR}(\boldsymbol{u}) \enspace ,$$

$$\textit{and furthermore,} \quad \mathrm{AR}(\boldsymbol{u}) \le \frac{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{u}'+\varepsilon,\boldsymbol{s})}{\inf_{\boldsymbol{s}\in\mathrm{E}_{2\varepsilon}(\boldsymbol{u}')} \mathrm{W}(\boldsymbol{u}'-\varepsilon,\boldsymbol{s})} \enspace .$$

[5]Some define the stability ratio as the reciprocal of $\mathrm{SR}(\boldsymbol{u})$.

Corollary 4.3. *Assume as in Cor. 3.3. We may then refine the lower bound, improving the supremum (i.e., maximally consonant) equilibrium value to an infimum (i.e., maximally dissonant), again so long as the denominator is positive, as*

$$\frac{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{u}'-\varepsilon,\boldsymbol{s})}{\inf_{\boldsymbol{s}\in\mathrm{E}(\boldsymbol{u}')} \mathrm{W}(\boldsymbol{u}'+\varepsilon,\boldsymbol{s}) + 2\gamma\varepsilon} \le \mathrm{AR}(\boldsymbol{u}) \enspace ,$$

$$\textit{and furthermore,} \quad \mathrm{SR}(\boldsymbol{u}) \le \frac{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{u}'+\varepsilon,\boldsymbol{s})}{\sup_{\boldsymbol{s}\in\mathrm{E}(\boldsymbol{u}')} \mathrm{W}(\boldsymbol{u}'-\varepsilon,\boldsymbol{s}) - 2\gamma\varepsilon} \enspace .$$

Next, paralleling the structure of our discussion of extreme equilibria, where we introduced Λ-stable outcomes as an alternative to extreme equilibria, we now introduce an alternative to the anarchy ratio, namely the anarchy gap.

Definition 4.4 (Anarchy and Stability Gaps). Given game Γ with utility function $\boldsymbol{u}$, the *anarchy gap* is defined as

$$\mathrm{AG}(\boldsymbol{u}) \doteq \underbrace{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{s})}_{\textsc{Cooperative}} - \underbrace{\inf_{\boldsymbol{s}\in\mathrm{E}(\boldsymbol{u})} \mathrm{W}(\boldsymbol{s})}_{\textsc{Non-Cooperative}} \enspace ,$$

and the *stability gap* is defined similarly.

Combining ideas, given $\Lambda \ge 0$, the Λ-*anarchy gap* is defined as the *Lagrangian relaxation* over the equilibrium set, i.e.,

$$\mathrm{AG}_\Lambda(\boldsymbol{u}) \doteq \underbrace{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{s})}_{\textsc{Cooperative}} - \underbrace{\inf_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{s}) + \Lambda\,\mathrm{Reg}(\boldsymbol{s};\boldsymbol{u})}_{\textsc{Non-Cooperative}} \enspace ,$$

and the Λ-*stability gap* is defined similarly.

If $\boldsymbol{u}$ has an equilibrium, then $\mathrm{AG}_\Lambda(\boldsymbol{u}) \ge 0$. However, as equilibria are not guaranteed to exist in games without mixed strategies, $\mathrm{AG}_\Lambda(\boldsymbol{u})$ can become very negative, as Λ tends toward infinity. Thus, we also define a relaxation of this gap: given $\Lambda \ge 0$, the Λ^*-*anarchy gap* is defined as

$$\mathrm{AG}^*_\Lambda(\boldsymbol{u}) \doteq \underbrace{\sup_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{s})}_{\textsc{Cooperative}} - \underbrace{\inf_{\boldsymbol{s}\in\boldsymbol{S}} \mathrm{W}(\boldsymbol{s}) + \Lambda\,\mathrm{Reg}^*(\boldsymbol{s};\boldsymbol{u})}_{\textsc{Non-Cooperative}} \enspace ,$$

where, as above, $\mathrm{Reg}^*(\boldsymbol{s};\boldsymbol{u})$ is the *excess regret*, defined as $\mathrm{Reg}^*(\boldsymbol{u},\boldsymbol{s}) \doteq \mathrm{Reg}(\boldsymbol{s};\boldsymbol{u}) - \inf_{\boldsymbol{s}'} \mathrm{Reg}(\boldsymbol{s}';\boldsymbol{u})$. As $\mathrm{Reg}^*(\boldsymbol{s};\boldsymbol{u}) \ge 0$ and $\sup_{x\in X} f(x) \ge \inf_{x\in X} f(x)$, it follows that $\mathrm{AG}^*_\Lambda(\boldsymbol{u}) \ge 0$.

These latter two measures of anarchy, parameterized by Λ, are Lipschitz properties, and thus are well behaved: i.e., they can be well approximated via a uniform approximation.

Theorem 4.5 (Approximating the Anarchy Gap). *Assume* $\mathrm{W} : \boldsymbol{S} \times (\boldsymbol{S} \to \mathbb{R}^{|P|}) \to \mathbb{R}$ *is a* λ_{W}*-Lipschitz property. It holds that* $\mathrm{AG}_\Lambda(\boldsymbol{u})$ *is* $2(\lambda_{\mathrm{W}}+\Lambda)$*-Lipschitz, which immediately implies*

$$\left|\mathrm{AG}_\Lambda(\boldsymbol{u}) - \mathrm{AG}_\Lambda(\boldsymbol{u}')\right| \le 2(\lambda_{\mathrm{W}}+\Lambda)\varepsilon \enspace .$$

It also holds that $\mathrm{AG}^*_\Lambda(\boldsymbol{u})$ *is* $4(\lambda_{\mathrm{W}}+\Lambda)$*-Lipschitz, which implies*

$$\left|\mathrm{AG}^*_\Lambda(\boldsymbol{u}) - \mathrm{AG}^*_\Lambda(\boldsymbol{u}')\right| \le 4(\lambda_{\mathrm{W}}+\Lambda)\varepsilon \enspace .$$

This theorem can be applied using one's preferred measure of welfare together with the corresponding Lipschitz constant. When an equilibrium is known to exist, we can use the regret property. If, however, an equilibrium is not known to exist (e.g., a pure strategy

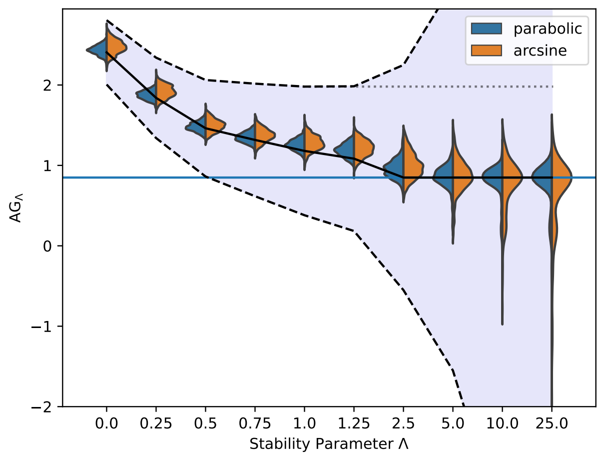

Figure 4: Approximations of Λ-anarchy gap via 0.2-uniform approximations of a randomly generated congestion game with 3 facilities, 3 players, and utility range $[-2, 2]$. 600 uniform approximations are randomly generated for both reduced bias (parabolic) and biased (arcsine) noise models.

equilibrium in a normal-form game), then we resort to using the excess regret property, and suffer a factor of two loss of accuracy.

All of the aforementioned results apply not only to the anarchy but to stability as well, for the corresponding definitions. Furthermore, analogs of these ratios and gaps can be defined for mixed games. As alluded to earlier, all our bounds hold, regardless of whether the game is pure or mixed.

The Λ-anarchy gap, by design, is a measure for which we can derive a more satisfying bound than the corresponding Λ-anarchy ratio. However, given an ε-uniform approximation of a game, this bound, like our bound on extreme Λ-stable outcomes, grows with Λ. Analogously, fixing the number of samples and letting Λ increase so as to subtract a value that is closer and to that of a pessimal equilibrium, yields a larger and larger confidence interval around the gap's estimate. Alternatively, fixing ε, as Λ grows, ε-accurate estimation of the gap requires more and more samples. We visualize this phenomenon in Figure 4 for utilitarian welfare.

Similar to Figure 1, this plot was created using GAMUT [18] to generate a congestion game $\boldsymbol{u}$ with 3 players, 3 facilities, and utility range $[-2, 2]$. We uniformly drew 600 normal-form games from the collection of ε-uniform approximations of this congestion game, by adding random noise $\varepsilon_{p,s} \leq \varepsilon \doteq 0.2$ to each individual utility, i.e., setting $\boldsymbol{u}'_p(\boldsymbol{s}) \doteq \boldsymbol{u}_p(\boldsymbol{s}) + \varepsilon_{p,s}$ for each $(p, \boldsymbol{s}) \in P \times \boldsymbol{S}$.

We experimented with two noise models: a mean-concentrated model and a biased model. In the mean-concentrated model, for each $(p, \boldsymbol{s})$, we draw $\varepsilon_{p,s}$ from the parabolic, i.e., $\beta(2, 2)$, distribution scaled to $[-\varepsilon, \varepsilon]$, which has (relatively small) standard deviation $\frac{\varepsilon}{4\sqrt{5}}$. The goal of this model is to have the utilities of the estimated game closer to the utilities of the ground-truth game (as one would expect via, for example, the central limit theorem) than worst-case tail bounds would mandate. In the biased model, for each $(p, \boldsymbol{s})$, we draw $\varepsilon_{p,s}$ from the arcsine distribution, i.e., $\beta(\frac{1}{2}, \frac{1}{2})$, scaled to range $[-\varepsilon, \varepsilon]$, which has the relatively large standard deviation of $\varepsilon/\sqrt{2}$. The goal of this model is to have the utilities of the estimated game farther away from the utilities of the ground-truth game, as more extreme games are more likely to exhibit large changes in $AG_\Lambda(\cdot)$.

Figure 4 shows several violin plots, one plot for each value of $\Lambda \in \{0, 0.25, 0.5, 0.75, 1.0, 1.25, 2.5, 5.0, 10.0, 25.0\}$. Notice that the selection of Λ's increases linearly in the first half, and exponentially in the second half, so as to cover a wide range of interesting behaviors. Each half of each violin plot shows the distribution of $AG_\Lambda(\boldsymbol{u}')$ across 600 0.2-uniform approximation $\boldsymbol{u}'$ of $\boldsymbol{u}$ for each of reduced bias and biased noise models. The plot also shows the upper and lower bound of Thm. 4.5, as dashed lines, as well as a dotted line representing an improved upper-bound on AG_Λ, as this quantity decreases monotonically in Λ (because the non-cooperative term is monotonic in Λ, and the cooperative term is constant). Finally, the true value of $AG_\Lambda(\boldsymbol{u})$ for each Λ is depicted via a solid black curve, and the true value of $AG(\boldsymbol{u})$ is represented by a blue line.

Our experiments reveal that our estimator is biased. Unsurprisingly, it appears less biased in the parabolic case than in the arcsine case, and the variance is also lower in the parabolic case. At $\Lambda = 25$, the empirical estimates are very noisy. This result is consistent with our observation that extreme equilibria cannot be well-estimated (Obs. 3.1). Even a very small change in utilities can effect whether a strategy profile is an equilibrium or not, and can thus lead to a large change in extreme equilibrium values.

5 CONCLUSION

Recent work [19] has argued against Nash, or any other static, equilibrium as the preferred solution concept, and in favor of alternatives that take into account the dynamics of learning agents, thereby deeming near-equilibrium behavior more acceptable than it has been in the past. Our results are yet another argument in this vein, with regard to optimal Nash equilibria, based not on computational complexity [9, 13], but rather on statistical estimation. In particular, there does not exist a finite m that is sufficient to estimate a welfare-optimizing Nash equilibrium well; on the other hand, estimating an extreme Λ-stable outcome for any finite Λ is a problem we solve in this paper.

Using the language of Λ-stable outcomes, the anarchy gap is defined as the difference between the maximally consonant 0-stable outcome and the maximally dissonant ∞-stable outcome. It is possible to imagine other gaps of interest, such as the maximally consonant Λ-stable outcome vs. the maximally dissonant Λ-stable outcome, for any finite Λ. The study of additional anarchy and stability gaps is left for future work.

So too is *empirical* mechanism design [25, 27] (EMD) using the ideas outlined in this paper. EMD is mechanism design with EGTA at its core, rather than traditional game-theoretic analysis. In the search for mechanisms, designers generally seek extreme solutions: e.g., extreme equilibria. As neither extreme equilibria nor the price of anarchy are well behaved, we suggest that EMD researchers optimize extreme Λ-stable outcomes or some version of an anarchy gap instead.

ACKNOWLEDGMENTS

This work was supported in part by NSF Grant CMMI-1761546.

REFERENCES

[1] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. 2004. The price of stability for network design with fair cost allocation. In *45th Annual IEEE Symposium on Foundations of Computer Science*. 295–304. https://doi.org/10.1109/FOCS.2004.68

[2] Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. 2020. Improved Algorithms for Learning Equilibria in Simulation-Based Games. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 79–87.

[3] Enrique Areyan Viqueira, Cyrus Cousins, Eli Upfal, and Amy Greenwald. 2019. Learning Equilibria of Simulation-Based Games. *arXiv preprint arXiv:1905.13379* (2019).

[4] Enrique Areyan Viqueira, Amy Greenwald, Cyrus Cousins, and Eli Upfal. 2019. Learning Simulation-Based Games from Data. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1778–1780.

[5] Pranjal Awasthi, Maria-Florina Balcan, Avrim Blum, Or Sheffet, and Santosh Vempala. 2010. On Nash equilibria of approximation-stable games. In *International Symposium on Algorithmic Game Theory*. Springer, 78–89.

[6] Gleb Beliakov, Tomasa Calvo, and Simon James. 2009. Some results on Lipschitz quasi-arithmetic means. In *IFSA-EUSFLAT 2009: Proceedings of the 2009 International Fuzzy Systems Association World Congress and 2009 European Society for Fuzzy Logic and Technology Conference*. EUSFLAT, 1370–1375.

[7] George Bennett. 1962. Probability inequalities for the sum of independent random variables. *J. Amer. Statist. Assoc.* 57, 297 (1962), 33–45.

[8] George Christodoulou and Elias Koutsoupias. 2005. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*. ACM, 67–73.

[9] Vincent Conitzer and Tuomas Sandholm. 2008. New complexity results about Nash equilibria. *Games and Economic Behavior* 63, 2 (2008), 621–641. https://doi.org/10.1016/j.geb.2008.02.015 Second World Congress of the Game Theory Society.

[10] Cyrus Cousins. 2021. An axiomatic theory of provably-fair welfare-centric machine learning. *Advances in Neural Information Processing Systems* 34 (2021), 16610–16621.

[11] Cyrus Cousins. 2022. Uncertainty and the Social Planner's Problem: Why Sample Complexity Matters. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 2004–2015.

[12] Cyrus Cousins. 2023. Revisiting Fair-PAC Learning and the Axioms of Cardinal Welfare. In *Artificial Intelligence and Statistics (AISTATS)*.

[13] Itzhak Gilboa and Eitan Zemel. 1989. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior* 1, 1 (March 1989), 80–93. https://doi.org/10.1016/0899-8256(89)90006-7

[14] Wassily Hoeffding. 1963. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association* 58, 301 (1963), 13–30.

[15] Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science*. Springer, 404–413.

[16] John Nash. 1953. Two-Person Cooperative Games. *Econometrica* 21, 1 (1953), 128–140. http://www.jstor.org/stable/1906951

[17] John F Nash. 1950. Equilibrium points in n-person games. *Proceedings of the national academy of sciences* (1950).

[18] Eugene Nudelman, Jennifer Wortman, Yoav Shoham, and Kevin Leyton-Brown. 2004. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2*. IEEE Computer Society, 880–887.

[19] Shayegan Omidshafiei, Christos Papadimitriou, Georgios Piliouras, Karl Tuyls, Mark Rowland, Jean-Baptiste Lespiau, Wojciech M Czarnecki, Marc Lanctot, Julien Perolat, and Rémi Munos. 2019. α-Rank: Multi-Agent Evaluation by Evolution. *Scientific Reports* 9 (2019).

[20] Tim Roughgarden, Vasilis Syrgkanis, and Éva Tardos. 2016. The Price of Anarchy in Auctions. *Journal of Artificial Intelligence Research* 59 (07 2016). https://doi.org/10.1613/jair.5272

[21] Andreas Schulz and Nicholás Stier-Moses. 2003. On the performance of user equilibria in traffic networks. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 86–87.

[22] Karl Tuyls, Julien Perolat, Marc Lanctot, Edward Hughes, Richard Everett, Joel Z Leibo, Csaba Szepesvári, and Thore Graepel. 2020. Bounds and dynamics for empirical game theoretic analysis. *Autonomous Agents and Multi-Agent Systems* 34, 1 (2020), 7.

[23] Karl Tuyls, Julien Perolat, Marc Lanctot, Joel Z Leibo, and Thore Graepel. 2018. A Generalised Method for Empirical Game Theoretic Analysis. *arXiv preprint arXiv:1803.06376* (2018).

[24] Leslie G Valiant. 1984. A theory of the learnable. *Commun. ACM* 27, 11 (1984), 1134–1142.

[25] Enrique Areyan Viqueira, Cyrus Cousins, Yasser Mohammed, and Amy Greenwald. 2019. Empirical mechanism design under uncertainty.

[26] Yevgeniy Vorobeychik. 2010. Probabilistic analysis of simulation-based games. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 20, 3 (2010), 16.

[27] Yevgeniy Vorobeychik, Christopher Kiekintveld, and Michael P Wellman. 2006. Empirical mechanism design: methods, with application to a supply-chain scenario. In *Proceedings of the 7th ACM conference on Electronic commerce*. ACM.

[28] Michael P. Wellman. 2006. Methods for Empirical Game-Theoretic Analysis (extended abstract). In *aaai06*. 1552–1556.

[29] John A Weymark. 1981. Generalized Gini inequality indices. *Mathematical Social Sciences* 1, 4 (1981), 409–430.

[30] Bryce Wiedenbeck, Ben-Alexander Cassell, and Michael P. Wellman. 2014. Bootstrap statistics for empirical games. 597–604.