Finding an Equilibrium in the Traveler's Dilemma with Fuzzy Weak Domination

Jesse Roberts
Computer Science Department
Vanderbilt University
Nashville, TN
Jesse.Roberts@vanderbilt.edu

***Abstract*—The traveler's dilemma is a well studied situation in the field of game theory in which the 2 players achieve very poor payoffs by each adhering to the Nash equilibrium strategy. Paradoxically, by deviating from the "rational" strategy, a significantly better payoff can be achieved, and it has been shown that humans overwhelmingly choose these "irrational" strategies. In this paper we (1) show that a single large assumption underlies the rational analysis which makes the Nash equilibrium rationally deducible and that (2) relaxing this assumption causes the deduction of the Nash equilibrium to fail. Therefore, we (3) introduce early work proposing an extension to game theory called fuzzy weak domination. Using fuzzy weak domination, we perform a rational analysis in which the opponent rationality assumption is regarded as uncertain and (4) show that the emergent equilibrium may more closely capture human behaviour.**

***Index Terms*—game theory, traveler's dilemma, weak domination, fuzzy weak domination**

I. Introduction

At the heart of game theory lies a pair of fundamental assumptions. First, it is assumed that each player is rational. That is, they will each make the best decision that they can. The second assumption prescribes that a decision is considered best if it maximizes the player's payoff. The rules of a given game and the rationale available to each player are considered common knowledge. So, given a space of possible strategies, a player will choose the strategy that is a best response to the strategy that they believe their opponent will pick.

If each player believes that their opponent is perfectly rational then an infinite sequence in which player A knows that player B knows that player A knows ... that player B knows some fact, is possible. However, a sequence of this type is so obviously intractable that it exists as a comedic trope [1]. A more recent sub-field called epistemic game theory attempts to formally reason about games given some set of beliefs held by the players. In this way, the players can be assumed to be rational and have common knowledge without requiring that they believe their opponent is rational.

In nature, groups of rational agents tend toward certain strategies in a game. These strategies are referred to as equilibria. While it is guaranteed that every game possesses at least one Nash equilibrium [2], a given game may have various other equilibria as well. Mathematically studying the rational analysis that underpins a given equilibria is important as analysis methods often generalize to other games, leading to the explanation or expectation of equilibria behaviour in these games as well.

In this paper we present early work in which we look at a set of results from an experiment involving a one-shot traveler's dilemma game. From an emergent equilibrium we argue that the human participants hold some level of uncertainty regarding their opponent's rationality. We then show that iterated elimination of weakly dominated strategies, which is used to find the Nash equilibrium in the traveler's dilemma, does not converge to the Nash equilibrium if players have non-zero uncertainty regarding opponent rationality. Finally, we present the first formulation of an extension to the idea of weak domination, referred to as fuzzy weak domination, which facilitates equilibrium analysis in the face of uncertainty regarding opponent rationality.

A. Traveler's Dilemma Paradox

We provide a short introduction to the traveler's dilemma (TD). For a more thorough discussion, see [3].

Suppose there are two people traveling back from vacation. Both of the travelers have purchased the same antique and have checked the antiques as luggage on the flight home. The airline breaks both antiques. The baggage claim team informs the two people of the broken antique and informs them that another passenger on the plane also had their identical antique broken.

The travelers are each told to give a value for the antique on the range $[2, 100]$, but they are warned that quoting a higher price than the other passenger will result in a penalty. Thus the airline has engaged the passengers in a 2-player game. If the two players provide the same quote, then they will each receive the amount quoted. However, given the quote from player A is Q_A and the quote from player B is Q_B, if $Q_A > Q_B$ then the payoff for player A will be $Q_B - 2$ and the payoff for player B will be $Q_B + 2$. The reciprocal statement is true if $Q_A < Q_B$.

This material is based upon work supported by the National Science Foundation under Grant No. 1521672.

978-1-6654-3886-5/21/$31.00 ©2021 IEEE

Definition 1. *A strategy is a Nash equilibrium iff no change in strategy can achieve a higher payoff assuming the opponent(s) does not change strategy.*

Partial and Total Ordering by Weak Domination

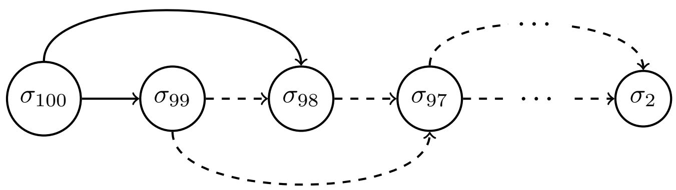

Fig. 1. An edge directed from vertex a to vertex b signifies a relative preference for vertex b. The solid edges are implied by $\sigma_{99} >_{wd} \sigma_{100}$ (Partial Ordering). While the dashed edges are implied by $\sigma_{98} >_{wd} \sigma_{99}$ etc. after weakly dominated strategies are eliminated due to zero uncertainty regarding opponent rationality (Extending to Total Ordering).

Definition 2. *For two strategies, σ_α is said to weakly dominate σ_β, that is $\sigma_\alpha \underset{wd}{>} \sigma_\beta$, iff for every opponent strategy, σ_α provides a payoff no worse than σ_β and there exists one or more scenarios in which σ_α provides a better payoff [4].*

The only Nash equilibrium strategy given in the format (Q_A, Q_B) for the TD is (σ_2, σ_2). It is obvious from definition 1 that (σ_2, σ_2) is a Nash equilibrium. However, the analysis that leads us to it is not obvious.

Consider that each player will likely begin to analyze strategies at a focal point [5]. In this case, we would expect the focal point to be either 2 or 100 with 100 being the most likely as it can provide a higher payoff. Each player considers σ_{100} and realizes that it is weakly dominated (definition 2) by σ_{99}, that is $\sigma_{99} \underset{wd}{>} \sigma_{100}$. In Fig.1, an edge directed from vertex a to vertex b signifies a relative preference for vertex b. So, there is an edge from σ_{100} to σ_{99} due to weak domination.

Each player then decides that it is impossible for their opponent to choose σ_{100} because they are **completely certain their opponent is rational** and in every situation, σ_{99} is as good or better than σ_{100}. So, the resulting set of available quotes after eliminating σ_{100} is $[2, 99]$. Each player then performs the same analysis and realizes that in the resulting set of available quotes, $\sigma_{98} \underset{wd}{>} \sigma_{99}$. Thus, σ_{99} is eliminated. This process is iterated until the only remaining option is σ_2. The preferences that are generated by iteratively eliminating the weakly dominated strategies are shown in Fig.1 as dashed lines.

II. Experimental Evidence

In [6], the authors ran a one-shot TD competition in which the competitors were drawn from the Game Theory Society. Each of the competitors submitted a strategy and the competitors were matched pairwise to every other competitor. The most successful strategy was σ_{97}. The mode of the strategies was σ_{100} with $N = 10$. The second most common was σ_{98} with $N = 9$. There were a total of 25 participants that used a strategy on the interval $[94, 99]$, 10 used $[100]$, 3 used $[2]$, and 7 on the interval $[4, 93]$. The authors of that paper theorize that there were 3 types of players involved. One type was an irrational player that played σ_{100}. The second was a rational player that played a best strategy given a belief about what others would play. The last was a type that either defaulted to the Nash equilibrium or started from a focal point of 2.

A. Non-zero Uncertainty

We focus on the rational type which settled into an equilibrium at σ_{98}. This study is interesting because the participants are people who have knowledge of game theory. Assuming that each player was attempting to maximize their payoff, it must be true that the players of this type were not certain that other players would not choose σ_{100} as their strategy. If the players were certain that opponents would not choose σ_{100}, they would have eliminated this as a strategy. In turn, this would have eliminated σ_{99}. And, if no player will choose σ_{99}, then σ_{98} will not maximize the payoff. In general we formalize this into proposition 1.

Proposition 1. *Let p be some rational player in the traveler's dilemma attempting to maximize their payoff. If p does not choose the Nash equilibrium strategy, σ_N, then it must be true that their belief regarding whether other players will choose σ_{100} has non-zero uncertainty. This is equivalent to having non-zero uncertainty regarding opponent rationality in general.*

Proof. If a player believes that their opponent will not choose σ_{100} with zero uncertainty due to weak domination, then it may be eliminated as a possibility. In the resulting space of possible strategies, $[2, 99]$, σ_{99} possesses all the same properties that led to the elimination of σ_{100}. Thus it is proved by induction. □

B. Ramifications for Analysis

From proposition 1, in order to account for the evidence in [6] we must consider that players may have a non-zero uncertainty regarding the rationality of opponents. However, if we allow uncertainty, the iterated elimination of weakly dominated strategies that we used to rationally deduce the Nash equilibrium fails. Consider, that if we can't eliminate σ_{100} then there is one possible opponent strategy on which σ_{98} does not provide a payoff equal to or better than σ_{99}. Therefore, σ_{98} does not weakly dominate σ_{99}. Visually, the effect is only the solid edges in Fig.1 can be deduced if there is any uncertainty regarding opponent rationality.

Even in the face of uncertainty, it seems intuitive that σ_{98} should be preferred over σ_{99} since the only possible scenario in which σ_{99} provides a better payoff is weakly dominated and therefore *unlikely*. To address this shortcoming, we will define a more general notion of weak domination, called fuzzy weak domination, that doesn't require certainty.

III. Fuzzy Logic

Here we will briefly introduce the core concepts of fuzzy logic.

Fuzzy logic was originally formulated in [7] to provide a method of reasoning in non-boolean contexts. As an example, consider a situation in which a recipe prescribes 2 minutes of boiling for large eggs and 1 minute for small eggs. By boolean logic, if an egg belongs to the set of small eggs it should be boiled for precisely 1 minute and 2 if an egg belongs to the set of large eggs. But what is the precise definition of large and small? Any value given to precisely identify the weight of a small egg and large egg will fail to cook the eggs properly unless the egg is the precisely prescribed weight.

In reality an egg may be somewhat large and somewhat small. That is, an egg can be considered to belong to both the set of large eggs and the set of small eggs with varying degrees of **membership** or **certainty**. Therefore, we can define a function with range $[0, 1]$ that fuzzifies the weight of the egg into the eggs membership in the fuzzy set of large eggs. We can likewise define a function that fuzzifies the weight of the egg into the eggs membership in the fuzzy set of small eggs. Then based on the certainty that an egg is in the fuzzy set of large eggs and the fuzzy set of small eggs an appropriate combination of the associated boiling times can be found.

Fuzzy logic necessarily redefines the logical operators common in boolean logic to work in an infinitely valued logic context. The result is that the boolean operator exists as a special case of the fuzzy operator. We quickly give the fuzzy operator definition for the NOT, AND, and OR operations.

The logical boolean NOT operator converts 1 to 0 and 0 to 1. The fuzzy **NOT** operation is defined as $1 - \mu$. Boolean AND operations return 1 if every operand in the operation is equal to 1. Fuzzy **AND** is equivalent to the min of the list of operands. Finally, boolean OR returns 1 if any of the operands are 1. The fuzzy **OR** operation returns the max of the list of operands.

IV. Fuzzy Weak Domination

In definition 3 we give a formal definition for fuzzy weak domination (FWD). To develop an intuitive understanding we will demonstrate how FWD allows us to reason about the ordering of strategies in the TD in the face of uncertainty.

Definition 3. *For two strategies, σ_α is said to fuzzy weak dominate σ_β with certainty $\zeta = \mu(s_\beta)$, that is $\sigma_\alpha \underset{fwd}{>} \sigma_\beta$, iff there exists one or more opponent strategies s.t. σ_α provides a better payoff than σ_β and the set of opponent strategies which provide a better payoff to σ_β must be a subset of the fuzzy set of all weakly dominated strategies with certainty $\zeta' > 0$. That is $s_\beta \in S_{fwd}$ with certainty $\zeta' > 0$.*

Let $\mu(s_\beta)$ be a membership function that transfers the membership in S_{fwd} from s_β to σ_β, where $s_\beta \in S_{fwd}$ with certainty ζ'. Further, if s_β contains more than one strategy, then the fuzzy membership of each element in the set is combined through an AND operation with a result equal to the minimum membership of any element in $s_\beta \in S_{FWD}$. Finally, for $y = \mu(x)$ it must be true that (1) $y \in [0, 1]$ $\forall x \in [0, 1]$, (2) $\exists x > 0$ s.t. $\mu(x) = 0$, and (3) if $x = 0$ then $y = 0$.

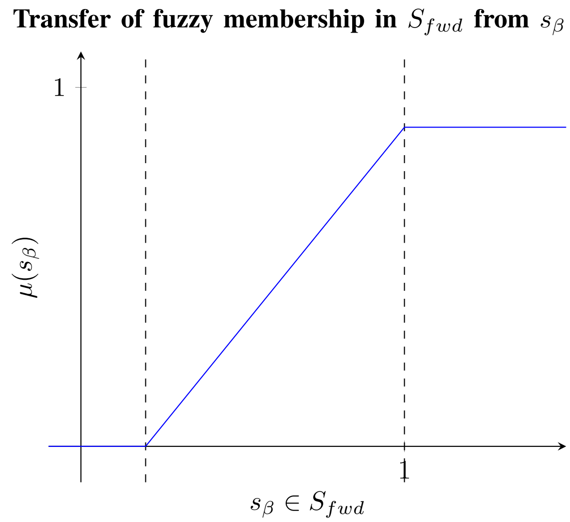

Fig. 2. A simple piecewise membership function for transference of fuzzy weak domination that satisfies the requirements in definition 3.

A. Example Execution of Partial Ordering by Iterated FWD

Initially the set of all FWD strategies contains only an empty set, $S_{fwd} = \{(\{\}, 1)\}$. Notice that S_{fwd} is a fuzzy set, so an entry in the set constitutes the value and membership pair. First, we find the set of opponent strategies for which σ_{99} provides a better payoff than σ_{100}. The resulting set is $\{\sigma_{100}, \sigma_{99}\}$. So, the first requirement in the definition evaluates to true. Next, we find the set of opponent strategies for which σ_{99} provides a worse payoff, s_β.

By the definition of FWD, we need to compute $\mu(s_\beta)$ and in this case $s_\beta = \{\}$. We can easily retrieve the membership associated with each element in s_β from the S_{fwd} and then apply the membership transference function, μ, to the minimum. A convenient μ is a simple piecewise linear function as shown in Fig.2 that satisfies the requirements in definition 3. The minimum membership of any element in $s_\beta = \{\}$ in S_{fwd} is 1 and $\mu(1) = 0.89$. We now update the set of FWD strategies to be $S_{fwd} = \{(\{\}, 1), (\sigma_{100}, 0.89)\}$. Now, rather than having absolute certainty that an opponent will not choose σ_{100} we say that σ_{100} is a member of the fuzzy weak dominated fuzzy set by with membership certainty equal to 0.89.

We iteratively apply this process and find the set of opponent strategies for which σ_{98} provides a better payoff than σ_{99}. This set is $\{\sigma_{99}, \sigma_{98}\}$. Next, we find the set of opponent strategies for which σ_{99} provides a worse payoff. This set is $s_\beta = \{\sigma_{100}\}$. The minimum membership of any element in the resulting s_β in S_{fwd} is 0.89 and $\mu(0.89) = 0.77$.

Continuing to apply iterative FWD yields a partial ordering of strategies that tends to uncertainty. After a small number of iterative steps the certainty goes completely to zero. In Fig.3 we show that μ in Fig.2 allows the strategies from σ_{100} to σ_{94} to be ordered based solely on fuzzy weak domination.

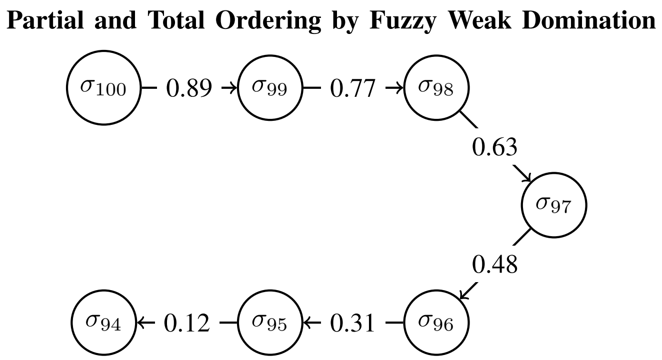

Fig. 3. An edge directed from vertex a to vertex b signifies a relative preference for vertex b.

B. Finding the Equilibrium in FWD

It is not initially obvious that the ordering in Fig.3 provides a specific prediction regarding the game's equilibrium. To see how the equilibrium emerges we consider the logical complement of each of the propositions.

For each member in the fuzzy set of fuzzy weak dominated strategies we can calculate our certainty that the member is not in the set. As an example, we can say that if $\sigma_{99} \underset{fwd}{>} \sigma_{100}$ with certainty 0.89 then it is also true that σ_{99} does not fuzzy weak dominate σ_{100} with certainty 0.11.

So, for a fuzzy partial ordering like that shown in Fig.3 there exists a natural equilibrium at the point of certainty inflexion. Notice that with certainty 0.63 $\sigma_{97} \underset{fwd}{>} \sigma_{98}$. However, we can say that $\sigma_{96} \underset{fwd}{\not>} \sigma_{97}$ with certainty 0.52. So, we are more certain that σ_{97} is not weakly dominated by σ_{96} than we are certain that it is weakly dominated. Therefore, a player attempting to maximize their payoff would not prefer σ_{96}.

It is not obvious if every partial ordering which results from iterated fuzzy weak domination possesses an inflexion point equilibrium.

C. Intuitive Understanding

The intuitive rationale that results from this analysis would be as follows. It is highly unlikely that the opponent will choose σ_{100} since σ_{99} is always as good or better. Since, σ_{100} is highly unlikely and σ_{98} is always as good or better than σ_{99} on every other strategy, it is unlikely that an opponent will choose σ_{99}. It still seems likely that σ_{97} is preferred because it provides a better payoff than σ_{98} on every strategy that is not highly unlikely or unlikely. At this point, certainty has fallen low enough that the player's belief has shifted such that the player has a higher certainty that σ_{96} does not fuzzy weak dominate σ_{97}. Thus the player chooses σ_{97}.

D. Comparing FWD to the Experimental Results

The results reported in [6] provided inspiration. Specifically, their work led us to consider that players may hold their opponents rationality as uncertain. That being said, it is notable that the equilibrium predicted by FWD with μ in Fig.2 is very close to the experimental result equilibrium (σ_{98}) as this was not engineered. We consider this to be affirmation (though not quite evidence) that FWD may accurately capture the rationale involved when humans engage in the TD.

It is more noteworthy that σ_{98} is the equilibrium in Fig.4 that corresponds to the largest interval when the x intercept of the general piecewise linear μ is swept on the domain $[0, 0.5]$. Therefore, if we assume that μ for an individual player may have an x intercept drawn randomly from the possibility space, then in general the most probable equilibrium is predicted to be σ_{98} by FWD. We consider this to be compelling evidence that FWD may accurately capture a specific type of rationale employed by the highest performing humans when engaged in the TD.

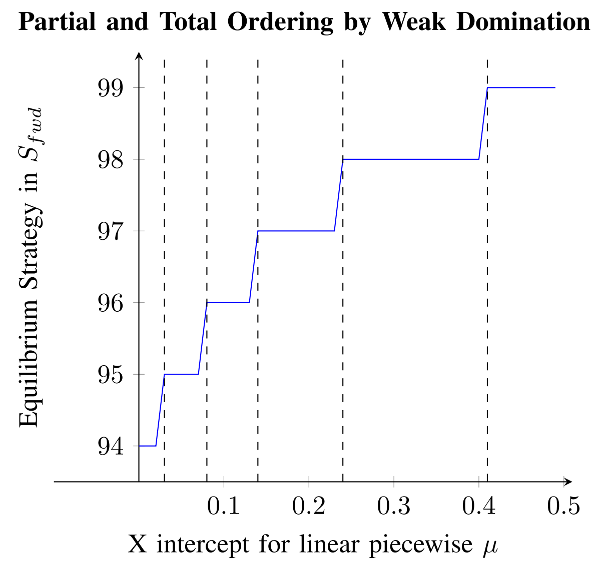

Fig. 4. Effect on the equilibrium predicted by FWD as a function of x intercept

V. Related Work

In [3] the author and original creator of the TD posits that it may be necessary to relax the rationality assumption in order to resolve the paradox. In this paper we do not relax the assumption that each player is rational. We instead relax the players' beliefs regarding the other's rationality.

The same author previously pointed out in [8] that fuzzy logic could be used to arrive at a better equilibrium in the TD. However, that discussion centered around treating the quote as a fuzzy value. Here we use fuzzy logic to deal with uncertainty regarding opponent rationality.

The experimental analysis in [6] cast the TD into a Bayesian game. They showed that if $P(\sigma_{100}) \geq 2\%$ then the Nash equilibrium of the Bayesian game is no longer located at σ_2. They continue to calculate the probability required to explain the results in the experiment given the Bayesian formulation. Probability and uncertainty are related but distinct ideas. Further, the intention of their analysis was to fit the data whereas FWD is formulated as a general analysis tool to find an expected equilibrium and natural extension of weak domination.

This is far from the first work to address the TD paradox. Neither is this the first paper to apply fuzzy logic to game theory. However, from our literature review this does seem to be the first paper to use fuzzy logic to address the inapplicability of iterated elimination of weakly dominated strategies when opponent rationality is not held certain in the TD.

VI. Discussion and Future Work

As already stated, the goal in defining FWD was not to fit the data in the experimental results of previous work. Rather, FWD was formulated as a fuzzy logic extension of weak domination to enable strategy ordering in the face of uncertainty regarding opponent rationality. Weak domination then exists as a special case of FWD in which μ is a unit step function written as $u(x - \tau)$ with $\tau > 0$. So, the fact that the experimental equilibrium from [6] emerges as the most probable FWD equilibrium suggests that FWD captures an important facet of rational thinking in the TD.

With that said, more testing in other games is needed to be able to evaluate if this rationale is applicable in a more general sense. It may also be found that FWD in the current formulation is incomplete. We intend to test this by changing the penalty involved in the TD game and comparing the effect on the predicted equilibrium against experimental results.

Another important point is that the choice of μ is non-trivial. More work is needed to evaluate the effect of other membership functions. An interesting and potentially promising extension would be the application of type 2 fuzzy logic so that one could formally reason when both the rationality of the opponent and the transference membership function are considered uncertain.

A. Application to Computational Sustainability

A method like FWD could potentially be applied to a wide range of games to identify likely equilibria. However, we specificially consider that FWD may prove useful in predicting the behaviour of opponents with uncertain rationality in games similar to the green security games defined in [9]. These green security games are important to the field of computational sustainability as they help to anticipate the actions of poachers and direct conservation efforts. FWD may be potentially well suited to this as green security games are derivative of Stackelberg security games, which possess open problems regarding scalability when faced with uncertainty [10].

VII. Conclusion

We have shown that by allowing a player to consider the opponent's rationality to be less than certain, iterated elimination of weakly dominated strategies does not provide a total ordering. In this scenario weak domination does not facilitate the deduction of a Nash equilibrium in the TD. We formulated an infinitely valued logic (fuzzy logic) extension of weak domination referred to as fuzzy weak domination. By iterated application of fuzzy weak domination we can generate a partial ordering. In the case of the TD, this partial ordering possesses an uncertainty inflexion point at which the complement of some partially ordered strategy's membership in the fuzzy set of fuzzy weak dominated strategies is greater than the membership itself. This inflexion point seems to be an equilibrium in the TD based on similarity to experimental results in [6].

Acknowledgment

The author thanks his advisor, Dr. Doug Fisher, and the National Science Foundation for support under Grant No. 1521672.

References

[1] *The Princess Bride*. MGM Home Entertainment, 2001.
[2] J. F. Nash *et al.*, "Equilibrium points in n-person games," *Proceedings of the national academy of sciences*, vol. 36, no. 1, pp. 48–49, 1950.
[3] K. Basu, "The traveler's dilemma," *Scientific American*, vol. 296, no. 6, pp. 90–95, 2007.
[4] A. Muthoo, M. J. Osborne, and A. Rubinstein, *A Course in Game Theory.*, 1996, vol. 63, no. 249.
[5] T. C. Schelling, "The strategy of conflict: Prospectus for a reorientation of game theory," *Journal of Conflict Resolution*, vol. 2, no. 3, pp. 203–264, 1958.
[6] T. Becker, M. Carter, and J. Naeve, "Experts Playing the Traveler's Dilemma," pp. 1–20, 2004. [Online]. Available: https://pdfs.semanticscholar.org/6c8d/b3a5fd2e87cd2a9bca9c31c896c441cf62b6.pdf
[7] L. A. Zadeh, "Information and control," *Fuzzy sets*, vol. 8, no. 3, pp. 338–353, 1965.
[8] K. Basu, "The traveler's dilemma: Paradoxes of rationality in game theory," *The American Economic Review*, vol. 84, no. 2, pp. 391–395, 1994.
[9] F. Fang, P. Stone, and M. Tambe, "When security games go green: Designing defender strategies to prevent poaching and illegal fishing." in *IJCAI*, 2015, pp. 2589–2595.
[10] A. Sinha, F. Fang, B. An, C. Kiekintveld, and M. Tambe, "Stackelberg security games: Looking beyond a decade of success." IJCAI, 2018.