

Games With Tolerant Players

Arpita Ghosh

Cornell University
Ithaca, NY 14853, USA
ag865@cornell.edu

Joseph Y. Halpern

Cornell University
Ithaca, NY 14853, USA
halpern@cs.cornell.edu

A notion of π -tolerant equilibrium is defined that takes into account that players have some tolerance regarding payoffs in a game. This solution concept generalizes Nash and refines ε -Nash equilibrium in a natural way. We show that π -tolerant equilibrium can explain cooperation in social dilemmas such as Prisoner's Dilemma and the Public Good game. We then examine the structure of *particularly cooperative π -tolerant equilibria*, where players are as cooperative as they can be, subject to their tolerances, in Prisoner's Dilemma. To the extent that cooperation is due to tolerance, these results provide guidance to a mechanism designer who has some control over the payoffs in a game, and suggest ways in which cooperation can be increased.

1 Introduction

People exhibit systematic deviations from the predictions of game theory. For example, they do not always act so as to maximize their expected utility in games such as Prisoner's Dilemma (to the extent that their utility is accurately characterized by the payoffs in the game). Many alternative models have been proposed to explain these deviations; the explanations range from players having *other-regarding* preferences, so that they prefer to avoid inequity, or prefer to maximize social welfare (see, e.g., [1, 5, 9]) to *quantal-response equilibrium*, which assumes that, with some probability, players make mistakes and do not play rationally [17].

The literature here is enormous; a complete bibliography would be almost as long as this paper. Nevertheless, we propose yet another approach to explaining what people do. Our explanation assumes that people have some *tolerance* for not getting the optimal payoff; the degree of tolerance is measured by how far from an optimal payoff they find acceptable. This idea is certainly not new: it is implicit in notions like ε -equilibrium and *satisficing* [20, 21], although the details are different. Moreover, it is clear that people do have some tolerance. There are many potential reasons for this. First, although we often identify payoffs and utilities, the payoffs in a game may not represent a player's true utility; a player may in fact be indifferent between receiving a payoff of a and $a - t$ if t is sufficiently small. (This is in the spirit of the "satisficing" view.) Or there may be a recommended strategy, and some overhead in switching (which again is not being captured in the game's payoffs). For whatever reason, it seems reasonable to assume that players may have some tolerance regarding payoffs. However, there is no reason to believe that all players have the same tolerance. We thus assume that there is a distribution over possible tolerances for each player, captured by a profile $\pi = (\pi_1, \dots, \pi_n)$ of distributions, where π_i is a distribution over the possible tolerances of player i .

Intuitively, we imagine that we have a large population of players who could be player i ; if we choose player i at random from this population, then with probability $\pi_i(t)$, she will have tolerance t . (Of course, in many applications, it is reasonable to assume that all players are chosen from the same pool, so that $\pi_1 = \dots = \pi_n$.) We can relate this to more traditional game-theoretic considerations by thinking of these tolerances as representing different *types* of player i ; that is, a type is associated with a tolerance.

There is some psychological evidence to support this viewpoint; specifically, there some evidence that whether someone uses satisficing-style behavior, and the extent to which it is used, is a personality trait, with a strong genetic component that endures over time [22]. In this setting, we define an equilibrium notion that we call π -tolerant equilibrium. A profile σ of possibly mixed strategies, one for each player, is a π -tolerant equilibrium if, roughly speaking, for each type t of player i (i.e., each possible tolerance that player i can have), we can assign a mixed strategy to type t in such a way that (1) each of the pure strategies in the mixture is a best response to σ_{-i} (i.e., what the other players are doing) and (2) σ_i represents the convex combination of what all the types of player i are doing. Intuitively, the other players don't know what type of player i they are facing; σ_i describes the aggregate behavior over all types of player i . We can show that a Nash equilibrium is a 0-tolerant equilibrium (i.e., if we take $\pi_1 = \dots = \pi_n$ to be the distribution that assigns probability 1 to players having tolerance 0); moreover, every Nash equilibrium is a π -tolerant equilibrium for all π . Similarly, if $\pi_1^\varepsilon = \dots = \pi_n^\varepsilon$ all assign probability 1 to players having tolerance ε , then a π^ε -tolerant equilibrium is an ε -Nash equilibrium. (The converse is not quite true; see Section 2.)

After defining π -tolerant equilibria in Section 2, in Section 3, we review the definition of social dilemmas, discuss the observed behavior in social dilemmas that we seek to explain, and show how tolerance can explain it.

Our interest in social dilemmas is only part of why we are interested in tolerance. We are also interested in taking advantage of tolerance when designing mechanisms. We illustrate the potential in Section 4 by investigating this issue in the context of Prisoner's Dilemma. Although Prisoner's Dilemma may seem to be a limited domain, it can model a range of two-player interactions with appropriate meanings ascribed to the actions of cooperating and defecting. Our analysis of Prisoner's Dilemma with tolerance isolates the factors that determine the equilibrium level of cooperation in the game, providing guidelines (to the extent to which tolerance is indeed the explanation for observed cooperation) for how a designer, who may be able to modify or control the payoffs from certain actions, can adjust them to achieve particular levels of cooperative behavior in equilibrium.

2 π -Tolerant Equilibrium

We consider normal-form games here. A *normal-form game* is a tuple $\Gamma = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$, where N is a finite set of *players*, which for convenience we take to be the set $\{1, \dots, n\}$, S_i is the set of pure strategies available to player i , which for convenience we also take to be finite, and u_i is i 's utility function. As usual, we take $S = S_1 \times \dots \times S_n$. A *mixed strategy* for player i is a distribution on S_i . Let Σ_i denote the mixed strategies for player i , and let $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$. Elements of Σ are called *mixed-strategy profiles*; given $\sigma \in \Sigma$, we denote by σ_i the i th component of the tuple σ , and by σ_{-i} the element of Σ_{-i} consisting of all but the i th component of σ . The utility function $u_i : S \rightarrow \mathbb{R}$; that is, u_i associates with each pure strategy profile a real number, which we can think of as i 's utility. We can extend u_i to Σ in the obvious way, by linearity.

We take T_i to be the set of possible tolerances for player i . Each element of T_i is a non-negative real number. For simplicity in this discussion, we take T_i to be finite, although the definitions that we are about to give go through with minimal change if T_i is infinite (typically, summations have to be changed to integrations). We identify a tolerance with a *type*; it can be viewed as private information about player i . Let π_i be a distribution on T_i , the set of possible types of i (under this viewpoint), and let $\pi = (\pi_1, \dots, \pi_n)$.

We want to define what it means for a mixed-strategy profile $(\sigma_1, \dots, \sigma_n)$ to be a π -tolerant equilibrium.

The intuition is that σ_i represents a population distribution. If σ_i puts probability p_i on the pure strategy s_i , then a fraction p_i of the population (of agents who could play the role of player i) plays s_i . Similarly, if π_i puts a probability p_i on a tolerance t , then a fraction p_i of the population of agents who could be player i has type t . Given our view that a mixed strategy for player i really represents a population of players each playing a pure strategy, in a π -tolerant equilibrium, we want all players of tolerance t to be playing a mixed strategy such that each strategy in the support is within a tolerance t of being a best response to what the other players are doing.¹

Definition 2.1 A pure strategy s_i for player i is *consistent with a tolerance t for player i and a mixed-strategy profile σ_{-i}* for the players other than i if s_i is a t -best response to σ_{-i} ; that is, if, for all strategies s'_i for player i ,

$$u_i(s'_i, \sigma_{-i}) \leq u_i(s_i, \sigma_{-i}) + t.$$

Definition 2.2 σ is a π -tolerant equilibrium if, for each player i , there is a mapping g_i from the set T_i of possible tolerances of player i to mixed strategies for player i such that the following conditions hold:

- E1. The support of $g_i(t)$ consists of only pure strategies that are consistent with t and σ_{-i} . (Intuitively, a player i of type t will play only strategies that are t -best responses to σ_{-i} .)
- E2. $\sum_t \pi_i(t) g_i(t) = \sigma_i$.

Note that if $(\sigma_1, \dots, \sigma_n)$ is a π -tolerant equilibrium, then there might not be any type of player i that plays strategy σ_i . Rather, σ_i describes the other players' perception of what a "random" instance of player i is doing. Thus, if player i has two possible types, say t and t' , where t occurs with probability $1/3$ and t' occurs with probability $2/3$, then E2 requires that $\sigma_i = \frac{1}{3}g_i(t) + \frac{2}{3}g_i(t')$.

Every Nash equilibrium is clearly a π -tolerant equilibrium for all π : For if σ is a Nash equilibrium, then each pure strategy in the support of σ_i is a best response to σ_{-i} , so must be consistent with t and σ_{-i} for all types t . Thus, if we take $g_i(t) = \sigma_i$ for all t , then E1 and E2 above are clearly satisfied. Moreover, the Nash equilibria are precisely the δ^0 -tolerant equilibria, where $\delta^0 = (\delta_1^0, \dots, \delta_n^0)$ and δ_i^0 puts probability 1 on type 0.

It is similarly easy to check that if $\delta^\varepsilon = (\delta_1^\varepsilon, \dots, \delta_n^\varepsilon)$, where δ_i^ε puts probability 1 on type ε , then every δ^ε -tolerant equilibrium is an ε -Nash equilibrium. The converse is *not* true, at least not the way that ε -Nash is typically defined (see, e.g., [18]).

For example, consider Prisoner's Dilemma. As is well known, defecting is the dominant strategy. Given $\varepsilon > 0$, there exists a $\delta > 0$ sufficiently small such that the mixed strategy $\delta C + (1 - \delta)D$ (cooperating with probability δ and defecting with probability $1 - \delta$) is an ε -best response no matter what the other player does; thus, both players using this strategy is an ε -Nash equilibrium. However, it is not a δ^ε -tolerant equilibrium if C is not an ε -best response. Interestingly, Goldberg and Papadimitriou [10] defined a (nonstandard) notion of ε -Nash equilibrium where all strategies in the support of a mixed strategy are required to be ε -best responses. This corresponds exactly to our notion of δ^ε -tolerant equilibrium.

Thus, π -tolerant equilibrium refines Nash equilibrium and ε -Nash equilibrium in an arguably natural way that allows for beliefs regarding agents' tolerance. We can also view it as a generalization of ε -Bayes-Nash equilibrium in Bayesian games. Recall that in a Bayesian game, each player i has a type in a set T_i . It is typically assumed that there is a (commonly known) distribution over $T = T_1 \cdots \times \cdots \times T_n$, and that a player's utility can depend on his type. The notion of ε -Bayes-Nash equilibrium in a Bayesian

¹We should stress that although we view a mixed strategy for player i as representing a population of players, each playing a pure strategy, nothing in the formal definitions requires this. There could equally well be a single player i playing a mixed strategy.

game a natural extension of ε -Nash equilibrium. If we take a player's type to be his tolerance, and take all types to agree on the utilities, then a π -tolerant equilibrium is a ε -Bayes-Nash equilibrium in the sense of the Goldberg-Papadimitriou definition, *provided that the ε can depend on the player's type*. That is, rather than having a uniform ε , we have a type-dependent ε . We believe that focusing on tolerance and its consequences gives more insight than thinking in terms of this nonstandard notion of Bayes-Nash equilibrium; that is what we do in the remainder of the paper.

We conclude this section by showing that greater tolerance leads to more equilibria. While this is intuitively clear, the proof (which can be found in the appendix) is surprisingly nontrivial. Given a distribution π_i , let F^{π_i} denote the corresponding cumulative distribution; that is, $F^{\pi_i}(t) = \sum_{t' \leq t} \pi_i(t')$. Say that π'_i *stochastically dominates* π_i if $F^{\pi'_i} \leq F^{\pi_i}$; that is, $F^{\pi'_i}(t) \leq F^{\pi_i}(t)$ for all t . Thus, the probability of getting greater than t with π'_i is at least as high as the probability of getting greater than t with π_i . Intuitively, π'_i stochastically dominates π_i if π'_i is the result of shifting π_i to the right. A profile $\pi' = (\pi'_1, \dots, \pi'_n)$ stochastically dominates $\pi = (\pi_1, \dots, \pi_n)$ if π'_i stochastically dominates π_i for all i .

Theorem 2.3 *If π' stochastically dominates π , then every π -tolerant equilibrium is a π' -tolerant equilibrium.*

Proof: See the appendix. ■

3 Social Dilemmas

Social dilemmas are situations in which there is a tension between the collective interest and individual interests: every individual has an incentive to deviate from the common good and act selfishly, but if everyone deviates, then they are all worse off. Following Capraro and Halpern [3], we formally define a social dilemma as a normal-form game with a unique Nash equilibrium σ^N and a unique welfare-maximizing profile s^W , both pure strategy profiles, such that each player's expected utility if s^W is played is higher than his utility if σ^N is played. While this is a quite restricted set of games, it includes many of the best-studied games in the game-theory literature.

We examine the same four games as Capraro and Halpern [3], and show that the experimentally observed regularities in these games can also be explained using tolerance.²

Prisoner's Dilemma. Two players can either cooperate (C) or defect (D). To relate our results to experimental results on Prisoner's Dilemma, we consider a subclass of Prisoner's Dilemma games where we think of cooperation as meaning that a player pays a cost $c > 0$ to give a benefit $b > c$ to the other player. If a player defects, he pays nothing and gives nothing. Thus, the payoff of (D, D) is $(0, 0)$, the payoff of (C, C) is $(b - c, b - c)$, and the payoffs of (D, C) and (C, D) are $(b, -c)$ and $(-c, b)$, respectively. The condition $b > c$ implies that (D, D) is the unique Nash equilibrium and (C, C) is the unique welfare-maximizing profile.

Traveler's Dilemma. Two travelers have identical luggage, which is damaged (in an identical way) by an airline. The airline offers to recompense them for their luggage. They may ask for any dollar amount between L and H (where L and H are both positive integers). There is only one catch. If they ask for the same amount, then that is what they will both receive. However, if they ask for different amounts—say one asks for m and the other for m' , with $m < m'$ —then whoever asks for m (the lower amount) will get $m + b$ (m and a bonus of b), while the other player gets $m - b$:

²The description of the games and observations is taken almost verbatim from Capraro and Halpern [3].

the lower amount and a penalty of b . (L, L) is thus the unique Nash equilibrium, while (H, H) maximizes social welfare, independent of b .

Public Goods game. $N \geq 2$ contributors are endowed with 1 dollar each; they must simultaneously decide how much, if anything, to contribute to a public pool. (The contributions must be in whole cent amounts.) The total contribution pot is then multiplied by a constant strictly between 1 and N , and then evenly redistributed among all players. So the payoff of player i is $u_i(x_1, \dots, x_N) = 1 - x_i + \rho(x_1 + \dots + x_N)$, where x_i denotes i 's contribution, and $\rho \in (\frac{1}{N}, 1)$ is the *marginal return*. (Thus, the pool is multiplied by ρN before being split evenly among all players.) Everyone contributing nothing to the pool is the unique Nash equilibrium, and everyone contributing their whole endowment to the pool is the unique welfare-maximizing profile.

Bertrand Competition. $N \geq 2$ firms compete to sell their identical product at a price between the “price floor” $L \geq 2$ and the “reservation value” H . (Again, we assume that H and L are integers, and all prices must be integers.) The firm that chooses the lowest price, say s , sells the product at that price, getting a payoff of s , while all other firms get a payoff of 0. If there are ties, then the sales are split equally among all firms that choose the lowest price. Now everyone choosing L is the unique Nash equilibrium, and everyone choosing H is the unique welfare-maximizing profile.³

From here on, we say that a player *cooperates* if he plays his part of the socially-welfare maximizing strategy profile and *defects* if he plays his part of the Nash equilibrium strategy profile. While Nash equilibrium predicts that people should always defect in social dilemmas, in practice, we see a great deal of cooperative behavior. But the cooperative behavior exhibits a great deal of regularity. Here are some regularities that have been observed (although it should be noted that in some cases the evidence is rather limited—see the discussion of Bertrand Competition at the end of this section):

- The degree of cooperation in the Prisoner's dilemma depends positively on the benefit of mutual cooperation and negatively on the cost of cooperation [4, 8, 19].
- The degree of cooperation in the Traveler's Dilemma depends negatively on the bonus/penalty [2].
- The degree of cooperation in the Public Goods game depends positively on the constant marginal return [11, 14].
- The degree of cooperation in the Public Goods game depends positively on the number of players [15, 23].
- The degree of cooperation in the Bertrand Competition depends negatively on the number of players [6].
- The degree of cooperation in the Bertrand Competition depends negatively on the price floor [7].

Of course, as mentioned in the introduction, there have been many attempts to explain the regularities that have been observed in social dilemmas. However, very few can actually explain all the regularities mentioned above. Indeed, the only approaches seem to be Charness and Rabin's [5] approach, which assumes that agents care about maximizing social welfare and the utility of the worst-off individual as well as their own utility, and the translucency approach introduced by Halpern and Pass [13] and adapted by Capraro and Halpern [3] to explain social dilemmas: roughly speaking, a player is translucent to the degree that he believes that, with some probability, other players will know what he is about to do.

³We require that $L \geq 2$ for otherwise we would not have a unique Nash equilibrium, a condition we imposed on Social Dilemmas. If $L = 1$ and $N = 2$, we get two Nash equilibria: $(2, 2)$ and $(1, 1)$; similarly, for $L = 0$, we also get multiple Nash equilibria, for all values of $N \geq 2$.

To show how tolerance can explain cooperation in social dilemmas, we first examine the relationship between tolerance and the parameters of the various social dilemmas we are considering. We consider two settings. In the first, we ask when it is consistent for a player i of type t to cooperate. In the second, we ask when it is rational for a player i of type t who believes (as assumed by Capraro and Halpern [3]) that each other player j is playing $\beta s_j^W + (1 - \beta)s_j^N$ to cooperate (i.e., i believes that each other player is cooperating with probability β , defecting with probability $(1 - \beta)$, and not putting positive probability on any other strategy). We write $\beta s_{-i}^W + (1 - \beta)s_{-i}^N$ for the corresponding mixed-strategy profile. Say that a player has type (t, β) in the second case; in the spirit of Definition 2.1, say that cooperation is *consistent with type (t, β) for player i* if for all strategies s'_i for player i ,

$$u_i(s'_i, \beta s_{-i}^W + (1 - \beta)s_{-i}^N) \leq u_i(s_i^W, \beta s_{-i}^W + (1 - \beta)s_{-i}^N) + t.$$

For a Prisoner's Dilemma of the form described in Section 3, consistency is independent of the strategy the other player is using (and hence independent of player's beliefs).

Proposition 3.1 *For the Prisoner's Dilemma of the form described in Section 3, cooperation is consistent for a player of type t and mixed strategy σ for the other player iff $t \geq c$.*

Proof: Switching from cooperating to defecting gives the player an additional payoff of c , independent of whether the other player is cooperating or defecting. Thus, cooperation is consistent if $t \geq c$. ■

We next consider the Traveler's Dilemma.

Proposition 3.2 *For the Traveler's Dilemma,*

- (a) *cooperation is consistent with t and a mixed strategy σ for the other player if $t \geq 2b - 1$;*
- (b) *there exists a strategy σ for the other player such that cooperation is consistent with t and σ iff $t \geq 2b - 1$;*
- (c) *cooperation is consistent with (t, β) iff $t \geq \max(\beta(b - 1), b - \beta(H - L))$.*

Proof: If player 2 plays m , then player 1's best response is to play $m - 1$ (or L if $m = L$). If $m < H$, then player 1 gets a payoff of $m - b$ if he cooperates (i.e., plays H), and could have gotten $m - 1 + b$ by making a best response (or L , in case $m = L$). Thus, he can gain at most $2b - 1$ by playing a best response. This proves part (a). If $m = H - 1$ and $H - 1 > L$, then cooperation is consistent iff $t \geq 2b - 1$; this proves part (b). Finally, if player 1 has type (t, β) , then he believes that 2 plays H with probability β and L with probability $1 - \beta$. Thus, player 1 believes his expected payoff from playing H is $\beta H + (1 - \beta)(L - b)$. The best response for player 1 is to play one of $H - 1$ or L . His payoff from playing $H - 1$ is $\beta(H + b - 1) + (1 - \beta)(L - b)$; his payoff from playing L is $\beta(L + b) + (1 - \beta)L$. Thus, cooperation is consistent for a player 1 of type (t, β) iff $t \geq \max(\beta(b - 1), b - \beta(H - L))$. ■

For the Public Goods game, consistency is again independent of the other players' strategies.

Proposition 3.3 *For the Public Goods game, cooperation is consistent for a player i of tolerance t and mixed strategy σ_{-i} for the other players iff $t \geq (1 - \rho)$.*

Proof: It is easy to see that, no matter what the other players do, defection (contributing 0) is the best response in this game, and a player gets a payoff that is $1 - \rho$ higher if he defects than if he cooperates. Thus, cooperation is consistent iff $t \geq (1 - \rho)$. ■

Proposition 3.4 *For Bertrand Competition with n players, cooperation for player i is consistent with (t, β) iff $t \geq \max(\beta^{n-1}(H-1), f(n)L) - \beta^{n-1}H/n$, where $f(n) = \sum_{k=0}^{n-1} \beta^k (1-\beta)^{n-1-k} \binom{n-1}{k} / (n-k)$.*

Proof: Consider a player of type (t, β) . If player i cooperates, he will get H/n if all the other players cooperate, which happens with probability β^{n-1} ; otherwise, he gets 0. Thus, his expected payoff from cooperation is $\beta^{n-1}H/n$. His best response, given his beliefs, is to play one of $H-1$ or L . If he plays $H-1$, then his payoff is $H-1$ if all the remaining players play H , which happens with probability β^{n-1} ; otherwise his payoff is 0. Thus, his expected payoff is $\beta^{n-1}(H-1)$. If he plays L , then his payoff if k players play H and $n-1-k$ play L is $L/(n-k)$; this event occurs with probability $\beta^k (1-\beta)^{n-1-k} \binom{n-1}{k}$. Thus, his expected payoff is $f(n)L$. It follows that cooperation is consistent with (t, β) if $t \geq \max(\beta^{n-1}(H-1), f(n)L) - \beta^{n-1}H/n$. ■

From here it is but three short steps to our desired result: First, observe that, up to now, we have looked at games in isolation. But now we want to compare tolerances in different games, with different settings of the relevant parameters. Intuitively, having a tolerance of 2 in Traveler's Dilemma where $L = 2$ and $H = 100$ should have a different meaning than it does in a version of Traveler's Dilemma where payoffs are multiplied by a factor of 10, so that $L = 20$ and $H = 1000$. Thus, when considering a family of related games, rather than considering absolute tolerances, it seems more appropriate to consider *relative tolerance*. There are many ways of defining a notion of relative tolerance. For our purposes, we take a player's relative tolerance to be an element of $[0, 1]$; player i 's actual tolerance in a game Γ is his relative tolerance multiplied by the payoff that player i gets if everyone cooperates in Γ . For example, since the payoff obtained by i if everyone cooperates in Traveler's Dilemma is H , then the actual tolerance of a player of type $(\tilde{t}, \beta)$ is $\tilde{t}H$. (Here and elsewhere, if we wish to emphasize that we are considering relative tolerance, we write $\tilde{t}$, reserving t for actual tolerance.) There are other ways we could define relative tolerance. For example, we could multiply by the difference between the payoff obtained if everyone cooperates and the payoff obtained if everyone defects, or multiply by the maximum possible social welfare. The exact choice does not affect our results.

Second, recall that the fact that cooperation is consistent with a given type does not mean that a player of that type will actually cooperate. We add an extra component to the type of a player to indicate whether the player will cooperate if it is consistent to do so, given his beliefs. We thus consider *relative types* of the form $(\tilde{t}, \beta, C)$ and $(\tilde{t}, \beta, D)$; such a type will cooperate in Traveler's Dilemma if $\tilde{t}H \geq \max(\beta(b-1), b - \beta(H-L))$ and the third component is C . Finally, we need to assume that there are a reasonable number of players of each type. Formally, we assume that the set of types of each player is infinite and that there is a distribution on relative types such that for all intervals (u, v) and (u', v') in $[0, 1]$, there is a positive probability of finding someone of relative type $(\tilde{t}, \beta, C)$ with $\tilde{t} \in (u, v)$ and $\beta \in (u', v')$. An analogous assumption is made by Capraro and Halpern [3].

With these assumptions, it follows from Propositions 3.1–3.4 that the regularities discussed in Section 3 hold.

- In the case of Prisoner's Dilemma, $b - c$ is the payoff obtained if everyone cooperates, so if $\tilde{t}$ is the relative tolerance, $\tilde{t}(b - c)$ is the actual tolerance. Thus, if a player's relative type is $(\tilde{t}, \beta)$, then cooperation is consistent if $\tilde{t}(b - c) \geq c$. Clearly, as b increases, there are strictly more relative types for which cooperation is consistent, so, by our assumptions, we should see more cooperation. Similarly, if c increases (keeping b fixed), there are fewer relative types for which cooperation is consistent, so we should see less cooperation.
- In the case of Traveler's Dilemma, as we have observed a relative type will cooperate if $\tilde{t}H \geq \max(\beta(b-1), b - \beta(H-L))$. Clearly, if b increases, then there will be fewer relative types for whom cooperation is consistent.

- In the Public Goods game, if everyone cooperates, the payoff to player i is $n\rho$. So it is consistent to cooperate if $\tilde{t}\rho n \geq (1 - \rho)$. Clearly, as n increases, we should see more cooperation, given our assumptions. Moreover, tolerance explains the increase of cooperation as the marginal return increases.
- Finally, in the Bertrand Competition, since the payoff if everyone cooperates is H/n , it is consistent to cooperate if $\tilde{t}H/n \geq \max(\beta^{n-1}(H - 1), f(n)L) - \beta^{n-1}H/n$, or equivalently, if

$$\tilde{t} \geq \max(n\beta^{n-1}(H - 1)/H, nf(n)L/H) - \beta^{n-1}.$$

Clearly, cooperation decreases if L increases. The effect of increasing n is more nuanced. For n large, β^{n-1} is essentially 0, as is $n\beta^{n-1}$; it can be shown that $f(n)$ is roughly $1/(1 - \beta)n$. Thus, if n is large, cooperation is consistent if $\tilde{t} > L/(1 - \beta)H$. What happens for small values of n is very much dependent on β , H , and L . The actual experiments on this topic [6] considered only 2, 3, and 4 players, with $L = 2$ and $H = 100$. For these values, we get the desired effect if β is sufficiently large ($\beta > .7$ suffices).

As we said earlier, of all the approaches to explaining social dilemmas in the literature, only Capraro and Halpern [3] and Charness and Rabin [5], can explain all these regularities; see Capraro and Halpern [3] for a detailed discussion and a comparison to other approaches. Of course, this leaves open the question of which approach is a better description of what people are doing. We suspect that translucency, care for others, and tolerance all influence behavior. We hope that further investigation of social dilemmas will reveal other regularities that can be used to compare our approach to others, and give us a more fine-grained understanding of what is going on.

4 Prisoner's Dilemma with Tolerance

We now take a closer look at the impact of tolerance on perhaps the best-studied social dilemma, Prisoner's Dilemma. The analysis suggests how thinking in terms of tolerance might help us design better mechanisms.

The general prisoners' dilemma (PD) game has payoffs (a_1, a_2) , (b_1, c_2) , (c_1, b_2) , and (d_1, d_2) corresponding to action profiles (C, C) , (C, D) , (D, C) , and (D, D) , respectively, with $c_i > a_i > d_i > b_i$. To analyze equilibrium outcomes in PD with tolerances, consider a player i , and suppose she believes that the other player j will play C with probability α_j . Her payoff from choosing action C and action D are, respectively,

$$u_C = a_j\alpha_j + (1 - \alpha_j)b_j; \quad u_D = \alpha_j c_j + (1 - \alpha_j)d_j.$$

Since D is a dominant strategy in Prisoner's Dilemma, $u_D > u_C$. Agent i is willing to play C if $u_D - u_C$ is within her tolerance, that is, if

$$t \geq \alpha_j(c_j - a_j) + (1 - \alpha_j)(d_j - b_j) \triangleq \alpha_j\Delta C_j + (1 - \alpha_j)\Delta D_j,$$

where ΔC_j is the gain to player i from defecting when the other player plays C , and similarly for ΔD_j . Taking F_i to denote the cumulative probability distribution on agent i 's tolerances, it follows that the probability that agent i has the tolerance required to allow cooperation is $1 - F_i(\alpha_j\Delta C_j + (1 - \alpha_j)\Delta D_j)$.

Note that the minimum tolerance at which i can cooperate, which depends upon the probability α_j with which j plays C , need *not* decrease with α_j : If the payoffs in the game are such that $\Delta C_j > \Delta D_j$ (the gain from defecting is larger when the other player is cooperating rather than defecting), then increasing α_j increases the tolerance i must have to be willing to cooperate.

Suppose that agents break ties in favor of cooperation, that is, if cooperating yields a payoff within an agent's tolerance, that agent will cooperate rather than defect. Call a π -tolerant equilibrium satisfying this condition a *particularly cooperative* (π -tolerant) equilibrium. Note that if some player with tolerance t cooperates in a perfectly cooperative equilibrium, then all players with tolerance t cooperate.

It follows from the discussion above that a particularly cooperative equilibrium is determined by a pair of mutually consistent probabilities of cooperation (α_i, α_j) satisfying

$$\begin{aligned}\alpha_i &= 1 - F_i(\alpha_i \Delta C_j + (1 - \alpha_i) \Delta D_j) \\ \alpha_j &= 1 - F_j(\alpha_j \Delta C_i + (1 - \alpha_j) \Delta D_i).\end{aligned}$$

Note that a particularly cooperative equilibrium may not exist in PD, although a π -tolerant equilibrium always does (since both agents defecting is a π -tolerant equilibrium, no matter what π is). For a simple example, suppose that $a_1 = a_2 = 3$, $b_1 = b_2 = -1$, $c_1 = c_2 = 5$, $d_1 = d_2 = 0$, and players are drawn from a population where everyone has a tolerance of 1.5. Now suppose that there is a perfectly cooperative equilibrium where a fraction α of the players cooperate. Thus, a player's expected payoff from cooperation is $3\alpha - (1 - \alpha) = 4\alpha - 1$; a player's expected payoff from defection is 5α . Thus, a player gains $\alpha + 1$ by defecting. If $0 \leq \alpha \leq .5$, then a player gains at most 1.5 by switching from cooperate to defect, so all players should cooperate in a perfectly cooperative equilibrium (i.e., α should be 1); on the other hand, if $\alpha > .5$, then $1 + \alpha > 1.5$, so all players should defect (so α should be 0).

In this example, we have a point mass of 1 on tolerance 1.5, so there is only one type of each player. This is inconsistent with the assumption in the previous section that the cumulative probability increases continuously. If we assume that the cumulative probability increases continuously then there is always a particularly cooperative equilibrium in PD. We provide an analysis here, making some symmetry assumptions for simplicity. Specifically, we assume (1) symmetry in payoffs: $a_1 = a_2, \dots, d_1 = d_2$; (2) symmetry in tolerance distributions: $F_1 = F_2 = F$; and (3) that F is continuous (so that there are infinitely many types). Under these assumptions, we show that a *symmetric* perfectly cooperative equilibrium (where $\alpha_1 = \alpha_2$) always exists; this is a solution α^* to

$$1 - \alpha = F(\alpha \Delta C + (1 - \alpha) \Delta D). \quad (1)$$

(Note that for prisoners' dilemma, $0 < F(\Delta C), F(\Delta D) \leq 1$ since $\Delta C = c - a$ and $\Delta D = d - b$ are positive.)

Theorem 4.1 (Equilibrium structure.) *Under our assumptions, a symmetric particularly cooperative equilibrium always exists, is a solution to (1), and has the following structure:*

- (a) *There is an equilibrium with $\alpha = 0$ (in which case (D, D) is necessarily played, so there is no cooperation) if and only if $F(\Delta D) = 1$.*
- (b) *(Uniqueness.) If $\Delta C > \Delta D$, there is a unique equilibrium; if $\Delta C \leq \Delta D$, multiple equilibria corresponding to different cooperation probabilities may exist.*

We omit the formal proofs; the results follow from applying the Intermediate Value Theorem using the continuity of F , and noting that the LHS in (1) is greater than the RHS at $\alpha = 0$ when $F(\Delta D) < 1$, and is smaller than the RHS at $\alpha = 1$. Uniqueness (and non-uniqueness, respectively) follows from the fact that the RHS is increasing (respectively, decreasing) in α when $\Delta D < \Delta C$ (respectively $\Delta C < \Delta D$). The idea is perhaps best illustrated by Figure 1, where the intersections correspond to equilibrium cooperation probabilities α^* . Note that these results depend critically on F , the cumulative distribution, being continuous.

The next result gives insight into how the probability of cooperation changes as we change various parameters. As we change the relevant parameters (the payoffs and the probabilities of tolerance) slightly




Figure 1: Equilibrium structure for PD.

in a continuous way, each particularly cooperative equilibrium “shifts” slightly also in a continuous way, so that we can talk about corresponding equilibria; we omit the formal definition here.

Theorem 4.2 *The equilibrium probability of cooperation α^* in corresponding particularly cooperative equilibria (a) decreases as ΔC and ΔD increase, and therefore (b) decreases as the payoffs c and d increase, and increases as a and b increase, and (c) “increases” in π , in that if π' stochastically dominates π , then the payoffs in a particularly cooperative π' -tolerant equilibrium are higher than in the corresponding π -tolerant equilibrium.*

These results follow easily by noting that for a fixed α , the RHS in (1) is increasing in both ΔC and ΔD ; the value of α at which the RHS equals the LHS therefore decreases when either of ΔC and ΔD increase. The monotonicity in π is similar: if π'_1 stochastically dominates π_1 , the value of α at the intersection is larger with π'_1 than with π_1 .

These results, in addition to providing testable predictions for how cooperation levels should behave when payoffs are varied in an experiment, also provide guidelines for a designer who may be able to manipulate some of the payoffs in the game (via either social or monetary means), by isolating the factors that influence the nature of equilibria and extent of equilibrium cooperation. First, the extent of cooperation in equilibrium depends on the marginal, rather than the absolute, payoffs: it is the *differences* ΔC and ΔD that determine equilibrium levels of cooperation, rather than any other function of the payoffs a, b, c, d .⁴

Second, perhaps surprisingly, which of ΔC or ΔD is larger makes a difference to the structure of equilibria, which also has implications for design. If, for instance, the designer prefers a game where there is a unique equilibrium with non-trivial cooperation, Theorem 4.2 suggests that the designer should manipulate payoffs so that ΔD , the marginal gain from defecting instead of cooperating when the other

⁴We must be a little careful here. To do comparative statics, we should consider relative tolerances, for the reasons discussed in Section 3. Changing the payoff parameters may well change the actual tolerance, while keeping the relative tolerance fixed. Parts (a) and (b) of Theorem 4.2 hold only for a fixed tolerance distribution.

player also defects, is smaller than ΔC , the marginal gain from defecting when the other player cooperates. (This might be achieved, for instance, by providing additional rewards—either extra compensation, or social rewards such as public acknowledgement, to a player who continues to cooperate despite defection by her partner, increasing the payoff b and therefore decreasing $d - b$.) On the other hand, if there is a means to “nudge” behavior towards a particular equilibrium when there are multiple equilibria, a designer might prefer to manipulate payoffs to fall in the $\Delta C \leq \Delta D$ regime and nudge behavior towards the equilibrium with the most cooperation (again, this could be achieved by imposing social or monetary penalties for defecting on a cooperating partner, decreasing t and thereby ΔC).

5 Conclusion

We have defined a notion of π -tolerant equilibrium, which takes into account that players have some tolerance regarding payoffs. This solution concept generalizes Nash and ε -Nash equilibrium in a natural way. We showed that this solution concept can explain cooperation in social dilemmas. Although we focused on social dilemmas, tolerance can also explain other well-known observations, such as the fact that people give some money to the other person in the *Dictator Game* [16] (where one person is given a certain amount of money, and can split it as he chooses between himself and someone else) and that people give intermediate amounts and reject small positive offers in the *Ultimatum Game* [12] (where one person can decide on how to split a certain amount of money, but the other person can reject the split, in which case both players get nothing).

We also examined the structure of particularly cooperative π -tolerant equilibria, where players are as cooperative as they can be, given their tolerances, in Prisoner’s Dilemma. To the extent that cooperation is due to tolerance, our results provide guidance to a mechanism designer who has some control over the payoffs in a game, and suggest ways that cooperation can be increased. Since many practical situations of interest can be modeled as Prisoner’s Dilemmas, these results may suggest how mechanism designers can take advantage of players’ tolerance in practice.

We believe that a study of convergence towards, and stability and robustness of, particularly cooperative equilibria in Prisoner’s Dilemma in an appropriate model for dynamics can potentially provide useful insights into emergence and sustainability of trust in online economies.

A Proof of Theorem 2.3

In this section, we prove Theorem 2.3. We repeat the statement of the theorem for the reader’s convenience

THEOREM 2.3. *If π' stochastically dominates π , then every π -tolerant equilibrium is a π' -tolerant equilibrium.*

Suppose that σ is a π -tolerant equilibrium, and π' stochastically dominates π for all i . We want to show that σ is a π' -tolerant equilibrium. Clearly, it suffices to consider the case where π'_i dominates π_i , and $\pi'_j = \pi_j$ for $j \neq i$.

Let the support of π_i be $\{t_1, \dots, t_n\}$, where $t_1 < \dots < t_n$, and let the support of π'_i be $\{t'_1, \dots, t'_m\}$, where $t'_1 < \dots < t'_m$. For convenience, define $t_0 = t'_0 = 0$. By assumption, there exists a mapping g_i such that $g_i(t)$ is a mixed strategy for each type t with support consisting of pure strategies consistent with t and σ_{-i} (this is E1) such that $\sum_{h=1}^n \pi_i(t_h) g_i(t_h) = \sigma_i$ (this is E2). We want to define a comparable function g'_i . In the remainder of the proof, for ease of exposition, we drop the subscript i (on g_i , g'_i , π_i , π'_i).

We start by defining $g'(t'_1)$. This has the benefit of giving the intuition for how to define g' in general. For ease of notation, we write F and F' rather than F^π and $F^{\pi'}$. Choose the smallest h such that $F(t_h) \geq F'(t'_1)$. Intuitively, type t'_1 should play $g(t_1)$ with probability $\pi(t_1)/\pi'(t'_1)$, $g(t_2)$ with probability $\pi(t_2)/\pi'(t'_1)$, ..., and $g(t_h)$ with probability $\pi(t_h)/\pi'(t'_1)$. This isn't quite right, since $\sum_{j=1}^h \pi(t_j) = F(t_h)$, and $F(t_h)$ may be strictly larger than $F'(t_1) = \pi'(t_1)$. Thus, we modify the probability that t'_1 plays $g(t_h)$ to $(\pi'(t_1) - F(t_{h-1}))/\pi'(t'_1)$. That is, $g(t_h)$ is played with whatever probability is left over after all the other strategies have been played. Thus, we take $g'_1(t'_1)$, the strategy played by type t'_1 , to be

$$\frac{\pi(t_1)}{\pi'(t'_1)}g(t_1) + \dots + \frac{\pi(t_{h-1})}{\pi'(t'_1)}g(t_{h-1}) + \frac{\pi'(t_1) - F(t_{h-1})}{\pi'(t'_1)}g(t_h). \quad (2)$$

We must show that E1 is satisfied by π' , so that any pure strategy in the support of any of $g(t_1), \dots, g(t_h)$ is consistent with t'_1 and σ_{-i} . Since consistency is monotonic in the tolerance, and by E1, all the strategies in the support of $g(t_j)$ are consistent with t_j and σ_{-i} , it suffices to show that $t'_1 \geq t_h$. Suppose, by way of contradiction, that $t'_1 < t_h$. By choice of t_h , if $t < t_h$, then $F(t) < F'(t'_1)$. Thus, $F(t'_1) < F'(t'_1)$. But this contradicts the assumption that F' stochastically dominates F . Thus, E1 is indeed satisfied by π' .

We now define $g'(t'_j)$ for $j \geq 2$. We first define auxiliary functions $\alpha : \{1, \dots, m\} \rightarrow \{1, \dots, n\}$ and $\beta : \{1, \dots, m\} \rightarrow [0, 1]$:

- $\alpha(j)$ is the least h such that $F(t_h) \geq F'(t'_j)$ (so the h in the definition of $g'(t'_1)$ above is just $\alpha(1)$).
- $\beta(j)$ is defined by induction on j . Let $\beta(1) = \pi'(t_1) - F(t_{\alpha(1)-1})$. Note that $\beta(1)$ was the quantity that occurred in the argument above that represented the amount of the probability mass $\pi'(t'_1)$ not allocated to $g(t_1), \dots, g(t_{\alpha(1)-1})$, which can thus be allocated to $g(t_{\alpha(1)})$. For $j \geq 2$, define

$$\beta(j) = \pi'(t_j) - \left(\pi(t_{\alpha(j-1)}) - \beta(j-1) + \sum_{h=\alpha(j-1)+1}^{\alpha(j)-1} \pi(t_h) \right).$$

Again, roughly speaking, $\beta(j)$ is the amount of the probability mass $\pi'(t'_j)$ not allocated to the strategies $g(t_{\alpha(j-1)}), \dots, g(t_{\alpha(j)})$.

We claim that, for $j \geq 2$, $t'_j \geq t_{\alpha(j)}$. The argument is essentially the same as that given above for the case that $j = 1$. Suppose not. Then, by choice of $t_{\alpha(j)}$, if $t < t_{\alpha(j)}$, then $F(t) < F'(t'_j)$. Thus, $F(t'_j) < F'(t'_j)$. But this contradicts the assumption that F' stochastically dominates F .

Finally, for $j > 1$, define

$$g'(t'_j) = \frac{1}{\pi'_1(t'_j)} \left[(\pi(t_{\alpha(j-1)}) - \beta(j-1))g(t_{\alpha(j-1)}) + \left(\sum_{h=\alpha(j-1)+1}^{\alpha(j)-1} \pi(t_h)g(t_h) \right) + \beta(j)g(t_{\alpha(j)}) \right]. \quad (3)$$

To see that this works, we need to check E1 and E2. For E1, note that the support of $g'(t'_j)$ consists of the strategies in the support of $g(t_{\alpha(j-1)}), \dots, g(t_{\alpha(j)})$. Since strategies in the support of $g(t_h)$ are all consistent with t_h and σ_{-i} , all the strategies in the support of $g'(t'_j)$ are clearly consistent with $t_{\alpha(j)}$ and σ_{-i} . We showed above that $t_{\alpha(j)} \leq t'_j$, so all these strategies are consistent with t'_j and σ_{-i} . Thus E1 holds.

For E2, note that it is immediate from (2) and the definition of $\alpha(1)$ and $\beta(1)$ that

$$\pi'_1(t'_1)g'_1(t'_1) = \sum_{h=1}^{\alpha(1)-1} \pi(t_h)g(t_h) + \beta(1)g(t_{\alpha(1)}). \quad (4)$$

It is immediate from (3) that for $j > 1$,

$$\pi'_1(t'_j)g'_1(t'_j) = (\pi(t_{\alpha(j-1)}) - \beta(j-1))g(t_{\alpha(j-1)}) + \left(\sum_{h=\alpha(j-1)+1}^{\alpha(j)-1} \pi(t_h)g(t_h) \right) + \beta(j)g(t_{\alpha(j)}). \quad (5)$$

It follows from (4) and (5) that

$$\sum_{h=1}^m \pi'(t'_h)g'(t'_h) = \sum_{h=1}^n \pi(t_h)g(t_h).$$

Since the latter sum is σ_1 by E2, we are done. ■

Acknowledgments: We thank Valerio Capraro for useful comments on an earlier draft of the paper. The work of Arpita Ghosh was supported in part by NSF grant CCF-1512964 and ONR grant N00014-15-1-2335 the work of Joe Halpern was supported in part by NSF grant CCF-1214844, AFOSR grant FA9550-12-1-0040, and ARO grants W911NF-14-1-0017 and W911NF-16-1-0397.

References

- [1] G. E. Bolton & A. Ockenfels (2000): *ERC: A theory of equity, reciprocity, and competition*. *American Economic Review* 90(1), pp. 166–193, doi:10.1257/aer.90.1.166.
- [2] M. Capra, J. K. Goeree, R. Gomez & C. A. Holt (1999): *Anomalous behavior in a traveler's dilemma*. *American Economic Review* 89(3), pp. 678–690, doi:10.1257/aer.89.3.678.
- [3] V. Capraro & J. Y. Halpern (2015): *Translucent players: explaining cooperative behavior in social dilemmas*. In: *Theoretical Aspects of Rationality and Knowledge: Proc. Fifteenth Conference (TARK 2015)*, pp. 87–94, doi:10.4204/EPTCS.215.9.
- [4] V. Capraro, J. J. Jordan & D. G. Rand (2014): *Heuristics guide the implementation of social preferences in one-shot Prisoner's Dilemma experiments*. *Scientific Reports* 4, doi:10.1038/srep06790.
- [5] G. Charness & M. Rabin (2002): *Understanding social preferences with simple tests*. *Quarterly Journal of Economics* 117(3), pp. 817–869, doi:10.1162/003355302760193904.
- [6] M. Dufwenberg & U. Gneezy (2000): *Price competition and market concentration: an experimental study*. *International Journal of Industrial Organization* 18, pp. 7–22, doi:10.1016/S0167-7187(99)00031-4.
- [7] M. Dufwenberg, U. Gneezy, J. K. Goeree & R. Nagel (2007): *Price floors and competition*. *Economic Theory* 33, pp. 211–224, doi:10.1007/s00199-006-0152-0.
- [8] C. Engel & L. Zhurakhovska (2016): *When is the risk of cooperation worth taking? The Prisoner's Dilemma as a game of multiple motives*. *Applied Economics Letters* 23(16), pp. 1157–1161, doi:10.1080/13504851.2016.1139672.
- [9] E. Fehr & K. Schmidt (1999): *A theory of fairness, competition, and cooperation*. *Quarterly Journal of Economics* 114(3), pp. 817–868, doi:10.1162/003355399556151.
- [10] P. W. Goldberg & C. H. Papadimitriou (2006): *Reducibility among equilibrium problems*. In: *Proc. 38th ACM Symposium on Theory of Computing*, pp. 61–70, doi:10.1145/1132516.1132526.
- [11] A. Gunthorsdottir, D. Houser & K. McCabe (2007): *Dispositions, history and contributions in public goods experiments*. *Journal of Economic Behavior and Organization* 62(2), pp. 304–315, doi:10.1016/j.jebo.2005.03.008.
- [12] W. Güth, R. Schmittberger & B. Schwarze (1982): *An experimental analysis of ultimatum bargaining*. *Journal of Economic Behavior and Organization* 3(4), pp. 367–388, doi:10.1016/0167-2681(82)90011-7.

- [13] J. Y. Halpern & R. Pass (2013): *Game theory with translucent players*. In: *Theoretical Aspects of Rationality and Knowledge: Proc. Fourteenth Conference (TARK 2013)*, pp. 216–221.
- [14] M. R. Isaac, J. M. Walker & S. Thomas (1984): *Divergent evidence on free riding: an experimental examination of possible explanations*. *Public Choice* 43(1), pp. 113–149, doi:10.1007/BF00140829.
- [15] M. R. Isaac, J. M. Walker & A. W. Williams (1994): *Group size and the voluntary provision of public goods*. *Journal of Public Economics* 54, pp. 1–36, doi:10.1016/0047-2727(94)90068-X.
- [16] D. Kahneman, J.L. Knetsch & R. H. Thaler (1986): *Fairness and the assumptions of economics*. *Journal of Business* 59(4), pp. S285–300, doi:10.1086/296367.
- [17] R. McKelvey & T. Palfrey (1995): *Quantal response equilibria for normal form games*. *Games and Economic Behavior* 10(1), pp. 6–38, doi:10.1006/game.1995.1023.
- [18] R. B. Myerson (1997): *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA.
- [19] A. Rapoport (1965): *Prisoner's Dilemma: A Study in Conflict and Cooperation*. University of Michigan Press, doi:10.3998/mpub.20269.
- [20] H. A. Simon (1955): *A behavioral model of rational choice*. *Quarterly Journal of Economics* 49, pp. 99–118, doi:10.2307/1884852.
- [21] H. A. Simon (1956): *Rational choice and the structure of the environment*. *Psychological Review* 63(2), pp. 129–138, doi:10.1037/h0042769.
- [22] I. Simonson & A. Sela (2011): *On the heritability of consumer Decision making: an exploratory approach for studying genetic effects on judgment and choice*. *Journal of Consumer Research* 37(6), pp. 951–966, doi:10.1086/657022.
- [23] J. Zelmer (2003): *Linear public goods experiments: A meta-analysis*. *Experimental Economics* 6, pp. 299–310, doi:10.1023/A:1026277420119.