



Interfaces with Other Disciplines

Limited-trust equilibria

Timothy Murray*, Jugal Garg, Rakesh Nagi

Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, United States



ARTICLE INFO

Article history:

Received 15 November 2019

Accepted 4 July 2020

Available online 10 July 2020

Keywords:

Game Theory

Limited-trust

Reciprocity

Long-term interactions

ABSTRACT

We propose and analyze a new type of equilibrium, in which limited-trust exists between players with long-term interactions. We assume heterogeneous interactions: players will engage in several games over an undetermined period of time with payoffs for each game drawn from a distribution. As such, players may not engage in the same game more than once. We define a Limited-Trust equilibrium to address these heterogeneous games, show its existence in all finite simultaneous games, and analyze it in general and in several common classes of games. We provide several interpretations of this equilibrium in leader-follower games. We then numerically compare the social utility generated from these equilibria in both simultaneous and leader-follower games to that generated by Nash and Stackelberg equilibria in the same games: when players display a similar level of trust δ , each sees an average gain of approximately δ in its utility each game over what it would achieve in traditional competitive/rational games, meaning for each game a player loses δ , there is another game it gains 3δ . Thus while players appear to play “non-rationally” by giving something up, they actually gain more and are each able to come out ahead of what they would have received if playing rationally as in a Nash equilibrium.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

One of the first things children learn is to “play nice” with others. In order to get ahead and be a functioning member of society, each individual must sometimes make choices which do not appear to benefit them in the short term. Even though these actions cost the individual, they make up for it in benefit to society; over time each individual will have these costs returned to it in the form of unexpected favors. Under the rationality considered in a Nash equilibrium (Nash (1950)), it makes sense to pay these costs and only violate the social rules when the cost is too great. This is particularly true if, after a certain amount of time, the rule violation will be forgiven or forgotten. To that end, a great deal of research has gone into the study of extensive form games in general, and repeated games in particular. These games frequently evince equilibrium behaviors which, when only considered for the individual stages rather than the extensive game, are not rational under the Nash definition. An explicit discussion of the work on these games will be presented in Section 1.1.

However, when there are no formal consequences to avoiding the costs of society, such as exile, why do individuals continue to incur these costs? For example, why will most people give up a seat on a bus to a stranger who is injured? In a Nash equilibrium,

in which only the utility of the individual making a decision is considered, the seat is never given up unless keeping it incurs some cost, such as damage to one's reputation. However, even if no one they know is present or will ever know of the decision, most people still give up the seat.

Perhaps the simplest answer is that the individual in the seat cannot know if their decision will ever make it back to others they interact with regularly and so they are simply risk averse. Another is that humans have some intrinsic degree of altruism. Evolutionary biology provides the best explanation of this in the form of Hamilton's rule for kin selection (Hamilton (1963, 1964a, 1964b)) which says that as humans are collections of genes, our genes seek to help any of the same genes present in other humans. To quote J. B. S. Haldane, “I would lay down my life for two brothers or eight cousins”. This idea that we lend aid to others because they are some proportion of ourselves has given rise to the concept of α -altruism, which will be discussed in Section 1.1.

In this paper we introduce a new concept which we refer to as a limited-trust equilibrium. In it, a player i attempts to maximize its long-term utility by trusting the other player(s) within a hard trust limit δ_i that it is willing to give up when the other player(s) will gain “significantly” more than it (they) would lose if player i were to play “rationally”. The player does this with the hope that the other player(s) will return the favor in a similar way, as well as form lasting partnerships and attract new ones through reputation. Given an opportunity, if an individual must choose between two agents of relatively equal capabilities to partner with

* Corresponding author.

E-mail addresses: tsmurra2@illinois.edu (T. Murray), jugal@illinois.edu (J. Garg), nagi@illinois.edu (R. Nagi).




Fig. 1. Hierarchy of equilibria. Intersection of all classes occurs in constant sum games.

then the individual would prefer to interact with the more trustworthy agent. We show through numerical trials that in two player games, when both players have a similar trust limit, $\delta_1 = \delta_2 > 0$, both players come out significantly ahead in the long term compared to if they had played solely to maximize their own utility: in 2-player numerical trials with $\delta_1 = \delta_2 = \delta$ we observe an average personal utility increase of δ for each player when δ was modest compared to the value of the variance in the utilities of randomly generated games.

The limited-trust equilibrium provides a new answer to the previous question of why someone would give up their seat on the bus to an individual who is injured: they do so to establish and contribute to a culture of “kindness”, which will increase the likelihood of someone giving them a seat in the event that they become ill or injured. This interpretation can be viewed as a person avoiding the consequences of the Broken Windows Theorem (Wilson & Kelling (1982)) which (loosely) states that evidence of erosion of one norm leads to further erosion of that and similar norms.

While it will be discussed more fully in Section 1.1, the idea of *non-rationality* within repeated games has been extensively studied. Therefore we pause briefly to distinguish this concept from other solution concepts which occur within repeated games: in such situations, the same games are played repeatedly and so players arrive at a best way to handle that single game over time using methods such as future discounting and trigger strategies. In limited-trust games, while players are assumed to be playing with each other over time, they are not assumed to play the exact same game continuously. In fact, they may never play the same game twice. Because of this, it is necessary that one-off games be analyzed individually, as each game may be independent of previous or later games played. This is something that other tools for repeated game analysis cannot do. If two players do interact again, the game will most likely be different as it is assumed to be drawn from some probability distribution.

The rest of this paper is organized as follows: In Section 1.1 we provide a more detailed discussion of previous work into extended form games as well as α -altruism. In Section 2 we fully detail the properties of a limited-trust equilibrium (LTE): we show that it is guaranteed to exist in finite n -player games, prove where it fits within the hierarchy of equilibrium concepts (see Fig. 1 for these results), and show that it results in higher net utility than Nash equilibria on several common games. Section 3 provides a mathematical program for LTE computation, and Section 4 discusses sev-

eral interpretations of limited-trust in the leader-follower setting. In Section 5 we present the results of numerical trials in both the simultaneous and leader-follower settings, in which we compare the highest value Nash equilibria to the highest value LTE's for randomly generated games, before moving to our final discussion of results and concluding remarks in Section 6.

1.1. Literature Review

Since the seminal work of Nash (1950) there has been a great interest in Game Theory and equilibrium concepts. In particular, many papers have noted that the strict definition of rationality adhered to by Nash equilibria, that it is a state where no player can unilaterally improve its own utility given the actions of other players, is frequently *not* observed in empirical trials. One circumstance in which this occurs is repeated games in which players engage in multiple rounds of play. Various *folk theorems* have been considered for these games which attempt to guarantee various measures of fairness in the equilibria; detailed analyses of these theorems and the conditions necessary for them to apply has been the subject of papers such as Aumann and Shapley (1994); Benoit and Krishna (1985); Fudenberg and Maskin (1986); Rubinstein (1979), and Rubinstein (1980). In the more applied sense, there has been a great deal of work aimed at developing rational definitions of trust for repeated games: papers such as Dasgupta (2000); Engle-Warnick and Slonim (2006b); Fudenberg and Maskin (1986) and Gibbons (2001) provide theoretical analysis of various games and trust strategies while papers such as Leopold-Wildburger, Schuetze, and Lafer (2002), Anderhub, Engelmann, and Gth (2002); Bapna, Qui, and Rice (2017); Berg, Dickhaut, and McCabe (1995); Engle-Warnick and Slonim (2004), and Engle-Warnick and Slonim (2006a) have focused on conducting empirical studies on several of these trust strategies, particularly in the context of reciprocity. In the business setting Morreale, Mitton, and Nigro (2019) experimentally tests the real options games approach put forward by lo Nigro, Morreale, Robba, and Roma (2013) for trust in strategic alliances. Meanwhile in the context of supply-chain relationships Chen and Wu (2010) empirically studies the formation of partnerships in the automotive industry, Du, Nie, Chu, and Yu (2014) derives a model for reciprocal-minded supplier-retailer relationships, and Niederhoff and Kouvelis (2016) and Choi and Messinger (2016) empirically show supply-chain relationships tend to be more “fair” over time than predicted in standard game theory. The recent survey Dal Bó and Frchette (2018) details many of these as well as other empirical studies, all of which on average show *non-Nash* behavior.

These trust papers, both theoretical and experimental, deal explicitly with repeated games or (as in Berg et al. (1995) and Leopold-Wildburger et al. (2002)) one-off games in extensive form (leader-follower). However, there is less work considering “non-rational” behavior in simultaneous one-off games. Most such work is done in the framework of α -altruism, as proposed by Ledyard (1994). In this concept, each player i has a perceived utility of $u_i^*(\sigma) = (1 - \alpha_i)u_i(\sigma) + \alpha_i u(\sigma)$ for $\alpha_i \in [0, 1]$ and thus takes the total social utility into account as part of its personal “utility”. This model is attractive for a number of reasons: it is supported by Hamilton’s kin-selection rule in evolutionary biology (Hamilton (1963, 1964a, 1964b)), it allows for easy equilibrium computation via Nash equilibria over perceived utilities, and it provides a broad model which can be adapted to virtually any form of game including simultaneous, extensive form, and repeated games. Chen, Keijzer, Kempe, and Schäfer (2014) provides a thorough analysis of this concept when applied to congestion, valid utility, and cost-sharing games, building on the analysis of Caragiannis, Kaklamanis, Kanellopoulos, Kyropoulou, and Papaioannou (2010) of this concept and extending the definition of (λ, μ) -smoothness put forth in an

earlier version of Roughgarden (2015) to α -altruistic games. However, this notion of altruism also has disadvantages, particularly from a modeling perspective. First, the game is scale invariant. This means that if player i would prefer not to collect € 1 so that player j can collect an extra € 2 given α_i , then it would prefer not to collect € 100 so that player j can collect an extra € 200 for the same α_i . Second, in games between a large number of players, the players are likely to become completely self-sacrificing to increase the total utility even for small $\alpha_i > 0$. To see this, consider a scenario in which for every unit of utility player i gives up, all other players receive some small amount of utility c , where $0 < c \ll 1$. As the number of players grows, player i will seek to drive its personal utility as low as possible so long as $\alpha_i > 0$.

In the next section we will propose a new concept of a limited-trust equilibrium which applies to a similarly broad class of games, but incorporates a hard trust-limit not present in α -altruism. Players behave in a manner which encourages reciprocity, provided it is not too expensive for them personally in terms of a hard limit on their current personal utility. They make this investment in reciprocation in order to increase their personal utilities in the long run or in expectation. This concept places a “budget” on what players spend toward encouraging reciprocity in any one game and thus eliminates both the tendency of players in large games to become self-sacrificing and the scale invariance which occur in α -altruistic games.

2. Limited-Trust Equilibrium

We now define a new concept of equilibrium in which players, while still selfish and concerned primarily with their own utility, exhibit a limited interest in the common good and contribute to it provided the cost is below some threshold. They do so in order to encourage other players to do the same in order to benefit in the long term. For comparison, we first review the definition of a mixed Nash equilibrium (MNE) over a finite game:

Definition 1 (Strategy Profile of a Finite Game). Given a finite n -player game in which each player i has m_i pure strategies, a valid strategy profile σ_i for player i is a probability distribution over the m_i pure strategies ($\sigma_i = \{p_1^i, p_2^i, \dots, p_{m_i}^i\}$, $p_j^i \geq 0$, $\sum_{j=1}^{m_i} p_j^i = 1$).

Definition 2 (Mixed Nash Equilibrium). Given an n -player game with strategy profiles $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ for each player where for a given player i , σ_{-i} is the set of strategies played by all other players, σ is a mixed Nash equilibrium (MNE) if and only if for any other valid strategy profiles σ'_i , $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$ for all $i \in [n]$, where $[n] = \{1, 2, \dots, n\}$ and $u_i(\sigma_i, \sigma_{-i})$ is the expected utility of the game for player i .

A related concept is the ε -approximate Nash equilibrium (ε -equilibrium) defined as follows:

Definition 3 (ε -Approximate Nash Equilibrium). For an n -player game with strategy profiles $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ for each player, σ is an ε -equilibrium if and only if for any other valid strategy profiles σ'_i , $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i}) - \varepsilon$ for all $i \in [n]$.

Definition 4 (Price of Anarchy). The Price of Anarchy (PoA) of a utility maximization game is the ratio of the value of the socially optimal solution, defined as the solution that maximizes the sum of the utilities of all players (net utility), to the value of the equilibrium with the lowest social utility.

Typically the equilibrium considered in the PoA is the Nash equilibrium; in this paper we will explicitly state which equilibrium is being considered when using the term.

Note that the set of Nash equilibria is merely the set of ε -approximate equilibria for $\varepsilon = 0$. It is also worth noting that the

conditions of an MNE can be defined in mathematical constraints. For an n -player utility maximization game, any strategy profile σ comprises an MNE if and only if it satisfies the following constraints:

$$u_i(\sigma'_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) \leq 0 \quad \forall \sigma'_i \in \Sigma_i, i \in [n]$$

where Σ_i is the set of valid strategy profiles for player i . We also define

$$\sigma_i^G(\sigma_{-i}) = \arg \max_{\sigma_i} u_i(\sigma_i, \sigma_{-i})$$

as the greedy best response of player i given σ_{-i} . We will abuse notation to let $\sigma_i^G \in \sigma_i^G(\sigma_{-i})$; while there may be multiple elements of $\sigma_i^G(\sigma_{-i})$, as it is a set-valued function, we will only be concerned with σ_i^G with regard to the value $u_i(\sigma_i^G, \sigma_{-i})$ which is equal for all elements of $\sigma_i^G(\sigma_{-i})$. We say that an ε -equilibrium $\sigma \in \Sigma$ is *well-supported* if and only if for every player i , $u_i(\sigma_i^G, \sigma_{-i}) - u_i(s_j^i, \sigma_{-i}) \leq \varepsilon$ for every pure strategy s_j^i which is played in σ_i with non-zero probability. Note that any MNE is a well-supported ε -equilibrium for all $\varepsilon \geq 0$.

Having covered our preliminary definitions, we now propose a new concept of equilibrium.

Definition 5 (Limited-Trust Equilibrium (LTE)). Consider a finite n -player maximization game with strategy profiles $\sigma \in \Sigma = \Sigma_1 \times \dots \times \Sigma_n$ and trust levels $\delta = (\delta_1, \dots, \delta_n)$ for each player i , where $\delta_i > 0$. σ is a limited-trust equilibrium if and only if $u_i(\sigma_i^G, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) \leq \delta_i$ and $u(\sigma_i, \sigma_{-i}) \geq u(\sigma'_i, \sigma_{-i})$ for any other valid strategy profiles $\sigma'_i \in \Sigma_i$ such that $u_i(\sigma_i^G, \sigma_{-i}) - u_i(\sigma'_i, \sigma_{-i}) \leq \delta_i$, where $u(\sigma) = \sum_{i=1}^n u_i(\sigma)$ is the net utility.

We will use $\text{LTE}(\delta)$ to refer to an LTE for players with trust levels $\delta = \{\delta_1, \dots, \delta_n\}$. This definition is equivalent to saying that the following two conditions are met:

1. Player i cannot alter its strategy profile to increase its payoff by more than δ_i . In other words, it is not giving up more than δ_i it could be making by changing its behavior to take advantage of other players' strategies.
2. Player i cannot alter its strategy profile to increase the net utility without decreasing its own utility so that it loses more than δ_i from its greedy best response. In other words, it cannot increase the net utility without violating its cost threshold δ_i .

As for where δ comes from, it can be viewed as the degree to which an individual is willing to invest in the future, meaning the cost they are willing to incur in order to benefit others and encourage them to reciprocate.

It is worth noting that we could equivalently write the net utility $u(\sigma_i, \sigma_{-i})$ as

$$u(\sigma_i, \sigma_{-i}) = \sum_{i=1}^n u_i(\sigma_i, \sigma_{-i}) = u_i(\sigma_i, \sigma_{-i}) + u_{-i}(\sigma_i, \sigma_{-i})$$

where $u_{-i}(\sigma_i, \sigma_{-i}) = \sum_{j \neq i} u_j(\sigma_j, \sigma_{-i})$. While there is no mathematical advantage in doing so, it helps to illustrate that if player i gives up δ_i by playing σ'_i rather than σ_i and the net utility increases by $x < \delta_i$, the $\delta_i - x$ value is not simply lost. Rather, if $u_i(\sigma'_i, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) = -\delta_i$ this means $u_{-i}(\sigma'_i, \sigma_{-i}) - u_{-i}(\sigma_i, \sigma_{-i}) = \delta_i + x$.

Because a limited-trust (LT) best response is concerned with two values $u(\sigma)$ and $u_i(\sigma)$ it makes sense to examine their relationship. In particular for a player i , if all other players are playing σ_{-i} then player i can easily determine the results of all of its pure strategies s_j^i in terms of $u(\sigma_i, \sigma_{-i})$ and $u_i(\sigma_i, \sigma_{-i})$. Because of the linearity of u and u_i with respect to s_j^i given a fixed σ_{-i} , any u_i combination within the convex hull of the pure strategies can be

achieved by player i . Therefore, player i can solve the following linear program LP1 to find its limited-trust best response $\sigma_i^*(\sigma_{-i})$:

$$\begin{aligned} \sigma_i^*(\sigma_{-i}) &= \arg \max_{\sigma_i \in \Sigma_i} u(\sigma_i, \sigma_{-i}) \\ \text{subject to} \quad & \delta_i \geq u_i(\sigma_i^G, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}). \end{aligned} \quad (\text{LP1})$$

When we take the limit $\delta_i \rightarrow \infty$ we find that player i becomes completely self-sacrificing for the net utility. Thus by careful selection of δ , players of any degree of trustworthiness from completely self-interested ($\delta \rightarrow 0$) to completely selfless ($\delta \rightarrow \infty$), may be modeled though players of the latter type may be quite uncommon.

We note that while the set $\sigma^G(\sigma) = \{\sigma_1^G(\sigma_{-1}), \sigma_2^G(\sigma_{-2}), \dots, \sigma_n^G(\sigma_{-n})\}$ is the set of greedy best responses to the current strategy set σ , $\sigma^G \in \sigma^G(\sigma)$ is not generally a Nash equilibrium. Instead, it is merely a set of greedy best responses to σ for each player. As such, the fact that σ_i^G is a component of player i 's limited-trust best response does not imply that a limited-trust equilibrium is dependent on a Nash equilibrium, merely that is dependent on greedy best responses. As Nash equilibria are also heavily dependent on greedy best responses, with $\sigma \in \sigma^G(\sigma)$ being a necessary and sufficient condition for σ to be a Nash equilibrium, this can be a subtle point. To further emphasize this distinction, we show in an example game in Table 4 that an LTE can exist independent of any Nash equilibrium.

Lemma 1. Given an $\text{LTE}(\delta)$ σ^* , if a constant c_j is added to all payoffs for player j , σ^* is still an $\text{LTE}(\delta)$.

The proof is included in Appendix A.1.

Although Lemma 1 demonstrates that an $\text{LTE}(\delta)$ is invariant under the addition of a constant c_j to all of player j 's payoffs, the same is not true for affine transformations. This is an intentional feature of the limited-trust concept: while a player may be willing to accept a loss of € 1 to ensure another player gains € 2, it is not willing to accept a loss of € 100 to ensure another player gains € 200, as would be required of an affine transformation of a game. However, for a given affine transformation $f(x) = ax + b$ the equilibria are invariant if δ is rescaled to $\delta|a|$, for $a \neq 0$.

Theorem 2. Every n -player finite game with trust levels $\delta = (\delta_1, \delta_2, \dots, \delta_n) > 0$ has an $\text{LTE}(\delta)$.

Proof. This proof will follow the same pattern as Nash's (Nash (1950)) proof for the existence of MNE in an n player game by making use of Kakutani's Fixed Point Theorem (Kakutani (1941)).

To begin, let $\sigma \in \Sigma$ be a set of strategy profiles for each player. Let $u_i(\sigma) = u_i(\sigma_i, \sigma_{-i})$ be the payoff player i derives from strategy profile σ_i given that all other players are playing σ_{-i} . Now we wish to define a new utility function

$$w_i(\sigma_i, \sigma_{-i}) = \begin{cases} u(\sigma_i, \sigma_{-i}) & u_i(\sigma_i^G, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) \leq \delta_i \\ -M & \text{otherwise,} \end{cases}$$

where M is a large positive constant. Because the game is finite, we can pick M greater than maximum of the absolute values of the socially optimal solution and the most socially harmful solution multiplied by n , and the $w_i(\sigma_i, \sigma_{-i})$ of any σ_i which violates $u_i(\sigma_i^G, \sigma_{-i}) - u_i(\sigma_i, \sigma_{-i}) \leq \delta_i$ is strictly less than $w_i(\sigma_i', \sigma_{-i})$ for some σ_i' which does not. Therefore, maximizing $w_i(\sigma_i, \sigma_{-i})$ is equivalent to maximizing $u_i(\sigma_i, \sigma_{-i})$ over the set of points which satisfy the maximum cost constraint. We can then say that $\sigma^* \in \Sigma$ is a $\text{LTE}(\delta)$ if and only if

$$w_i(\sigma_i^*, \sigma_{-i}^*) \geq w_i(\sigma_i, \sigma_{-i}^*) \quad \forall \sigma_i \in \Sigma_i, \forall i \in \{1, 2, \dots, n\},$$

which means that σ^* is a $\text{LTE}(\delta)$ if and only if $\sigma_i^* \in B_i(\sigma_{-i}^*)$ for all i , where $B_i(\sigma_{-i}^*)$ is the set of best responses (with respect to w_i) for player i given that the other players are playing σ_{-i}^* . If we define

$B(\sigma) = B_1(\sigma_{-1}) \times B_2(\sigma_{-2}) \times \dots \times B_n(\sigma_{-n})$ then finding an $\text{LTE}(\delta)$ is equivalent to finding $\sigma \in B(\sigma)$. Therefore, we must show the existence of a fixed point.

We now use Kakutani's Fixed Point Theorem to show such a fixed point exists. The theorem states that given a nonempty finite dimensional Euclidean space A and $f: A \rightarrow A$ a set-valued correspondence with $x \in A \rightarrow f(x) \subseteq A$, a fixed point is guaranteed to exist if the following conditions hold:

1. A is a compact and convex set.
2. $f(x)$ is nonempty for all $x \in A$.
3. $f(x)$ is convex for all $x \in A$.
4. $f(x)$ has a closed graph: if $\{x^k, y^k\} \rightarrow \{x, y\}$ with $y^k \in f(x^k)$ then $y \in f(x)$.

In this case we have $A = \Sigma$, $f(\sigma) = B(\sigma)$. We now wish to show that all conditions hold.

1. Σ is a compact and convex set: trivial, as Σ is the Cartesian product of simplices Σ_i .
2. $B(\sigma)$ is nonempty for all $\sigma \in \Sigma$: $B_i(\sigma_{-i}) = \arg \max_{\sigma_i \in \Sigma_i} w_i(\sigma_i, \sigma_{-i})$ and so must be nonempty for each i . Therefore $B(\sigma)$ is nonempty for all $\sigma \in \Sigma$.
3. $B(\sigma)$ is convex for all $\sigma \in \Sigma$: It suffices to show that $B_i(\sigma_{-i})$ is convex for all i . We first note that any points $x, y \in B_i(\sigma_{-i})$ must provide equal net utility $u(x, \sigma_{-i}) = u(y, \sigma_{-i})$ and must also provide i with a personal utility at most δ_i less than the greedy best response. Without loss of generality, assume $u_i(x, \sigma_{-i}) \geq u_i(y, \sigma_{-i})$. Then for any convex combination $z = \lambda x + (1 - \lambda)y$ where $\lambda \in [0, 1]$, the linearity of u and u_i implies that $u(x, \sigma_{-i}) = u(z, \sigma_{-i}) = u(y, \sigma_{-i})$ and $u_i(x, \sigma_{-i}) \geq u_i(z, \sigma_{-i}) \geq u_i(y, \sigma_{-i})$ which means $z \in B_i(\sigma_{-i})$.
4. $B(\sigma)$ has a closed graph: While the previous three conditions were shown to hold using the same arguments as in the proof of existence for Nash equilibria, the use of a non-continuous function w_i introduces several complications to showing that $B(\sigma)$ has a closed graph. We will show this by contradiction. Suppose that $B(\sigma)$ does not have a closed graph. Then there exists a sequence $(\sigma^k, \hat{\sigma}^k) \rightarrow (\sigma, \hat{\sigma})$ such that $\hat{\sigma}^k \in B(\sigma^k)$, but $\hat{\sigma} \notin B(\sigma)$, meaning that $\hat{\sigma}_i \notin B_i(\sigma_{-i})$ for some i . Then there is some $\sigma_i' \in B_i(\sigma_{-i})$ such that

$$w_i(\sigma_i', \sigma_{-i}) > w_i(\hat{\sigma}_i, \sigma_{-i}),$$

which means that

$$u(\sigma_i', \sigma_{-i}) > u(\hat{\sigma}_i, \sigma_{-i}).$$

By continuity of u_i and u , we have that for k sufficiently large $u(\sigma_i', \sigma_{-i}^k) > u(\hat{\sigma}_i^k, \sigma_{-i}^k)$. Because $\sigma_i' \notin B_i(\sigma_{-i}^k)$, we have $u_i(\sigma_i^{Gk}, \sigma_{-i}^k) - u_i(\sigma_i', \sigma_{-i}^k) > \delta_i$ where $\sigma_i^{Gk} \in \sigma_i^G(\sigma_{-i}^k)$ as otherwise this would contradict the assumption that $\hat{\sigma}_i^k \in B_i(\sigma_{-i}^k)$.

Suppose that $u_i(\sigma_i^{Gk}, \sigma_{-i}^k) - u_i(\hat{\sigma}_i^k, \sigma_{-i}^k) < \delta_i$. Then by the linearity of u_i and u there is a convex combination s of σ_i' and $\hat{\sigma}_i^k$ such that $u_i(\sigma_i^{Gk}, \sigma_{-i}^k) - u_i(s, \sigma_{-i}^k) \leq \delta_i$ and $u(s, \sigma_{-i}^k) > u(\hat{\sigma}_i^k, \sigma_{-i}^k)$, which implies $w_i(s, \sigma_{-i}^k) > w_i(\hat{\sigma}_i^k, \sigma_{-i}^k)$ and therefore contradicts the assumption that $\hat{\sigma}_i^k \in B_i(\sigma_{-i}^k)$.

Now suppose that $u_i(\sigma_i^{Gk}, \sigma_{-i}^k) - u_i(\hat{\sigma}_i^k, \sigma_{-i}^k) = \delta_i$. In order for σ_i' to become a strategy in $B_i(\sigma_{-i})$, it must be that $u_i(\sigma_i^G, \sigma_{-i}) - u_i(\sigma_i', \sigma_{-i}) = \delta_i$, as if it became less than or equal to δ_i for sufficiently high k , then it would contradict $\hat{\sigma}_i^k \in B_i(\sigma_{-i}^k)$. Similarly, $u_i(\sigma_i^G, \sigma_{-i}) - u_i(\hat{\sigma}_i, \sigma_{-i}) = \delta_i$ as if it is less than δ_i , then for sufficiently high k we would have $u_i(\sigma_i^{Gk}, \sigma_{-i}^k) - u_i(\hat{\sigma}_i^k, \sigma_{-i}^k) < \delta_i$ which we have already seen leads to a contradiction. However, this means that there is a strategy $\sigma_i'' = \lambda \sigma_i^G + (1 - \lambda) \sigma_i'$ for some value of $\lambda \in [0, 1]$ which has $u(\sigma_i'', \sigma_{-i}) > u(\hat{\sigma}_i, \sigma_{-i})$ and $u_i(\sigma_i'', \sigma_{-i}) > u_i(\hat{\sigma}_i, \sigma_{-i})$, due to the assumption that $\delta_i > 0$

Table 1
Example game for LTE.

		Player 2	
		β_1	β_2
Player 1	α_1	4,0	5,5
	α_2	5,1	6,2

and the linearity of u and u_i . This means that $u(\sigma_i'', \sigma_{-i}^k) > u(\hat{\sigma}_i^k, \sigma_{-i}^k)$ and $u_i(\sigma_i'', \sigma_{-i}^k) > u_i(\hat{\sigma}_i^k, \sigma_{-i}^k)$ for sufficiently high k . Given $u_i(\sigma_i^{Gk}, \sigma_{-i}^k) - u_i(\hat{\sigma}_i^k, \sigma_{-i}^k) \leq \delta_i$, this means $w_i(\sigma_i'', \sigma_{-i}^k) > w_i(\hat{\sigma}_i^k, \sigma_{-i}^k)$ which contradicts the assumption that $\hat{\sigma}_i^k \in B(\sigma_{-i}^k)$. Therefore $B(\sigma)$ must have a closed graph.

Therefore, Kakutani's theorem implies the existence of a $\sigma^* \in \Sigma$ such that $\sigma^* \in B(\sigma^*)$, which proves the existence of an $\text{LTE}(\delta)$. $\square$

Having established the guaranteed existence of an LTE for $\delta > 0$ we now want to compare it to a Nash equilibrium on a simple example, given in Table 1. This game has exactly one Nash equilibrium, at $\sigma = \{[0, 1], [0, 1]\}$ with pure strategies α_2, β_2 being played. Now consider the LTE with $\delta = \{0.5, 0.5\}$. LP1 shows that for player 2, playing the pure strategy β_2 ($\sigma_2 = [0, 1]$) is still the best choice, regardless of σ_1 . The same is not true for player 1: given $\sigma_2 = [0, 1]$, solving LP1 gives the first player's unique best response as $\sigma_1 = [0.5, 0.5]$. The net utility of the $\text{LTE}(0.5)$ for the game is then $u_1(\{[0.5, 0.5], [0, 1]\}) + u_2(\{[0.5, 0.5], [0, 1]\}) = 9$, compared to the net utility of 8 which occurs in the Nash equilibrium.

The equilibrium in Table 1 highlights an important fact about the limited-trust best response, that there may not be a pure strategy best response. This is at odds with the greedy best response where there is always a pure strategy best response. This implies that there does not appear to be a straightforward transformation of a limited-trust game into a Nash game.

Next, we wish to consider where the LTE fits within the hierarchy of standard solution concepts within game theory.

Theorem 3. For any finite n -player game G , the set of $\text{LTE}(\delta)$ is a subset of the set of ε -equilibria of G , where $\varepsilon \geq \max_i \delta_i$.

Proof. Consider that in any $\text{LTE}(\delta)$, no player can improve its own payoff by more than δ_i by definition of an LTE. Therefore such an LTE is also an ε -equilibrium for $\varepsilon \geq \max_i \delta_i$. $\square$

While each limited trust equilibrium is also a $\max_i \delta_i = \varepsilon$ -equilibrium, the converse is not true, even when $\delta_i = \delta_j$ for all $\forall i, j \in [n]$. This is because of the additional constraint on an $\text{LTE}(\delta)$ that no player i be able to improve the total utility without decreasing its own utility below the δ_i level. Further, although the set of $\text{LTE}(\delta)$'s is a subset of ε -equilibria as described in Theorem 3, they are important because they represent a state in which each player is contributing to the net utility as much as they are able within their limits, not merely a state where each player has decided it is not worth the effort (or in the case of irrational-valued Nash equilibria it is realistically infeasible) to change from their current strategy to the optimal strategy, particularly if the current strategy is pure.

In general, we say that an $\text{LTE}(\delta)$ σ is well-supported if it is a well-supported ε -equilibrium for $\varepsilon = \max_i \delta_i$. Although any $\text{LTE}(\delta)$ is an ε -equilibrium for ε as previously specified, it need not be a well-supported ε -equilibrium. The 2-player game in Table 2 demonstrates this for $\delta_1 = \delta_2 = 0.5$. From player 2's limited-trust perspective, β_2 is a best response to any σ_1 as it offers both better personal and better net utility. Player 1's limited-trust best response to β_2 is to play α_i with probability 0.5, for $i = \{1, 2\}$. Given that player 2 will only play β_2 , the only $\text{LTE}([.5, .5])$ is given by

Table 2
Game with non-overlapping $\text{LTE}([0.5, 0.5])$ and well-supported 0.5-equilibria.

		Player 2	
		β_1	β_2
Player 1	α_1	2,4	3,5
	α_2	3,2	4,3




Fig. 2. Best response curves for game in Table 3, $\delta_1 = \delta_2 = 0.1$.

Table 3
Game in which $\text{LTE}(0.1) < \text{MNE}$.

		Player 2	
		β_1	β_2
Player 1	α_1	1,2	5,0
	α_2	1,1,5,1	4,5

$\{[0.5, 0.5], [0, 1]\}$. This is not a well-supported 0.5-equilibrium: $\sigma_1^G = \alpha_2$ and player 1 is playing α_1 with nonzero probability, despite the fact that $u_i(\sigma_i^G, \beta_2) - u_i(\alpha_1, \beta_2) = 1 > 0.5$. As this is the only $\text{LTE}([0.5, 0.5])$, this game also shows that the set of $\text{LTE}(\delta)$ may be entirely disjoint from the set of well-supported ε -equilibria for a game.

Despite the fact that players in a limited-trust game all attempt to improve the net utility, it is possible for the highest value $\text{LTE}(\delta)$ (the $\text{LTE}(\delta)$ which provides the highest net utility) to produce lower net utility than the highest value Nash equilibrium for a game. We show this with an example game, given in Table 3, which we consider with $\delta_1 = \delta_2 = 0.1$. There is a pure Nash equilibrium (PNE) which occurs for α_2, β_1 . As β_1 is a strongly dominant strategy for player 2, and α_2 is player 1's best response to it, this is the only Nash equilibrium. If we consider the limited-trust best response curves in Fig. 2, we see that there is only one place the curves intersect and hence there is one $\text{LTE}(0.1)$. Using optimization program MP1 from Section 3, we find this is at approximately $\sigma_1 = \{p_1^1, p_2^1\} = \{0.204, 0.796\}$, $\sigma_2 = \{p_1^2, p_2^2\} = \{.795, .205\}$ which has a total utility value of approximately 6.089, which is less than 6.2, the utility generated by the pure Nash equilibrium at $\sigma_1 = \{0, 1\}$, $\sigma_2 = \{1, 0\}$.

While it is non-intuitive that the value of the best LTE can be less than a Nash equilibrium, given that each player is willing to give something up in order to help its fellow players, we do see analogues of this in the day-to-day social interactions which the concept of limited-trust emulates. Consider two cars reaching an intersection across from each other. Both need to turn left and the

Table 4
Game with more LTE(δ) than Nash equilibria.

		Player 2	
		β_1	β_2
Player 1	α_1	3,3	2,3.1
	α_2	3.1,2	5,5

intersection is too narrow for both to go at once. Rather than attempting to go through first, one driver tries to wave the other through, only to realize that the other driver is doing the same. Both drivers start to move, then stop as they realize the other is moving as well. This then repeats back and forth until one driver loses their patience (δ is reached) and makes it clear they are going. Meanwhile the whole interaction slowed down both drivers more than if one had simply made this decision when they both arrived at the intersection.

Although it is possible to find games in which there is a Nash equilibrium better than any LTE, we will see in Section 5 that it rarely occurs, particularly as δ increases; it is more common to find games in which there are more LTE(δ) than Nash equilibria and some of them are worse. Table 4 provides an example of this, where (α_2, β_2) is both an LTE(δ) and a pure Nash equilibrium. However, for $\delta_1, \delta_2 \geq 0.1$, (α_1, β_1) is also an LTE(δ), independent of a Nash equilibrium.

Further, the occurrence of less optimal solutions due to co-operation is not unique to limited-trust games: Chen et al. (2014) shows that while normal cost-sharing games have a PoA of n for n players, cost-sharing games in which all players have a uniform level of α -altruism have a PoA of $\frac{n}{1-\alpha}$, becoming unboundedly inefficient for fully altruistic players. The remainder of this section will be spent considering limited-trust versions of several standard games, and we will see that these inefficiencies do not apply to them.

2.1. LTE(δ) in Common Games

In this section we examine the behavior and value of LTE(δ) in several common classes of games.

Theorem 4. For any constant sum game, $\{\varepsilon\text{-equilibria}\} \subseteq \{\text{LTE}(\delta)\}$ where $\varepsilon = \min_i \delta_i$.

Proof. First, note that in a constant sum game the total utility is equal for all $\sigma \in \Sigma$. Therefore, any strategy σ played by player i maximizes the total utility. Thus, player i 's best response to any σ_{-i} is any strategy which makes sure it receives at most δ_i less than its maximum personal utility. This is exactly the definition of a best response under ε -equilibrium conditions for $\delta_i = \varepsilon$, and so the set of LTE contains the set of ε -equilibria for a constant sum game where $\varepsilon = \min_i \delta_i$. $\square$

Note that Theorem 3 states the LTE set is a subset of the ε -equilibria set for $\varepsilon = \max_i \delta_i$, so if $\delta_i = \delta_j = \delta$ for all $i \neq j$ then Theorem 4 implies the set of ε -equilibria is equal to the set of LTE(δ) for $\varepsilon = \delta$.

Next we consider the public goods game from experimental economics. In it, n players each receive an amount of money, m_i , and must decide how much to contribute to the public good. Any contributed money is multiplied by a factor of c such that $1 < c < n$, then divided evenly among all players. Therefore, if player i contributes x_i to the public good, it will receive back $\frac{cx_i}{n} < x_i$ of its investment, plus $\frac{c}{n}$ of the other players' investments. The only Nash equilibrium for this game is for all players to contribute $x_i = 0$, as any contribution lowers player i 's payoff regardless of contributions made by other players. The PoA is therefore c .




Fig. 3. Hotelling Game with $n = 3$ players.

Theorem 5. In a public goods game with δ , the limited-trust PoA is

$$\frac{c \sum_{i=1}^n m_i}{\sum_{i=1}^n m_i + (c-1) \min\{\frac{n}{n-c} \delta_i, m_i\}} \leq c.$$

Proof. Consider the contribution player i should make: the social utility strictly increases with i 's contribution x_i , therefore player i would like to contribute as much as possible. i is willing to lose at most δ_i and regardless of the value of x_j for $j \neq i$, if player i contributes x_i then it loses $\frac{n-c}{n} x_i$ it could be making. Therefore, player i contributes $x_i = \min\{\frac{n}{n-c} \delta_i, m_i\}$. The total amount contributed is $\sum_{i=1}^n x_i$, and the total uncontributed utility is $\sum_{i=1}^n m_i - x_i$, which means that the total utility generated is $\sum_{i=1}^n m_i + (c-1) \min\{\frac{n}{n-c} \delta_i, m_i\}$. This is the unique LTE(δ) for the public goods game. The socially optimal result occurs when all players contribute m_i and there is a total utility of $c \sum_{i=1}^n m_i$, so the limited-trust PoA is $\frac{c \sum_{i=1}^n m_i}{\sum_{i=1}^n m_i + (c-1) \min\{\frac{n}{n-c} \delta_i, m_i\}}$ that is at most c . $\square$

Our next consideration is the Hotelling game, which does not generally have a pure Nash equilibrium for any number of players n . We will consider the simplest form of the game, in which each player has a continuous strategy space $[0, 1]$ and all players have symmetric payoffs, meaning that for any two players i, j and all other strategies σ_{-ij} fixed, if i and j switched strategies they would also switch utilities ($u_i(\sigma_i = x_1, \sigma_j = x_2, \sigma_{-ij}) = u_j(\sigma_i = x_2, \sigma_j = x_1, \sigma_{-ij})$ for all $x_1, x_2 \in [0, 1]$). Given strategies $\{\sigma_1, \sigma_2, \dots, \sigma_n\}$, if we assume without loss of generality that $0 \leq \sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_n \leq 1$ then for $\sigma_{i-1} < \sigma_i < \sigma_{i+1}$ $u_i(\sigma) = \frac{\sigma_{i+1} - \sigma_{i-1}}{2}$. If there is a set of k strategies $\sigma_i = \sigma_{i+1} = \dots = \sigma_{i+k-1}$, then $u_j(\sigma) = \frac{\sigma_{i+k} - \sigma_{i-1}}{2k}, \forall i \leq j \leq i+k-1$. Additionally, for the purposes of computing σ_1 and σ_n , let " $\sigma_0'' = -\sigma_1$ " and " $\sigma_{n+1}'' = 1 + \sigma_n$ ".

This simple form of the Hotelling game can be viewed as each player claiming a space on the interval $[0, 1]$, with each player attempting to maximize the portion of the interval which is closer to them than all other players. Fig. 3 provides an example of this for a 3-player Hotelling game, which does not have a pure Nash equilibrium.

Theorem 6. The 3-player Hotelling game possesses a pure LTE(δ) for $\delta_i \geq \frac{1}{10}$ for $i \in \{1, 2, 3\}$.

Proof. This will be a proof by example, showing that $\sigma = \{\frac{3}{10}, \frac{1}{2}, \frac{7}{10}\}$ is an LTE(δ) for $\delta_i \geq \frac{1}{10}$ for $i \in \{1, 2, 3\}$. We begin by noting that the Hotelling game is constant-sum, so any strategy produces the same net utility. We first consider whether player 1 is at equilibrium. Observe that $u_1(\sigma) = \frac{2}{5}$, so player 1 is at equilibrium provided there is not some σ_1' such that $u_1(\sigma_1', \sigma_{-1}) > \frac{1}{2}$. As $\sigma_2 = \frac{1}{2}$, there cannot be: $\sigma_1' \in [0, \frac{1}{2})$, $(\frac{1}{2}, 1]$ will result in utility strictly less than $\frac{1}{2}$, and $\sigma_1' = \frac{1}{2}$ will result in the same utility as player 2 receives, which can be at most $\frac{1}{2}$ as the net utility for the game is 1. Therefore player 1 is at equilibrium, and similarly player 3 is at equilibrium as its position is symmetric to that of player 1.

Table 5
Example Prisoner's Dilemma Game.

		Player 2	
		β_1	β_2
Player 1	α_1	d_1, d_1	$d_2, 0$
	α_2	$0, d_2$	c, c

This leaves player 2. $u_2(\sigma) = \frac{1}{5}$, so it is at equilibrium if there is no σ'_2 such that $u_2(\sigma'_2, \sigma_{-2}) > \frac{3}{10}$. For $\sigma'_2 \in [0, \frac{3}{10})$, $(\frac{3}{10}, \frac{7}{10})$, $(\frac{7}{10}, 1]$, $u_2(\sigma'_2, \sigma_{-2})$ is $< \frac{3}{10}$, $\frac{1}{5}$, and $< \frac{3}{10}$ respectively. For $\sigma'_2 = \frac{3}{10}$, $u_2(\sigma'_2, \sigma_{-2}) = \frac{1}{4} < \frac{3}{10}$ as well, so player 2 is at equilibrium. $\square$

Before moving on, it is worthwhile to note that although [Theorem 2](#) only implies the existence of the $LTE(\delta)$ in finite normal form games, players in both the public goods game and Hotelling game possess a continuous rather than finite strategy set. This helps to highlight that while non-finite games are outside the scope of this paper, many classes of these games are also likely to possess limited-trust equilibria.

Finally, we consider the 2×2 prisoner's dilemma, though we will focus on the utility maximization version rather than the cost minimization version. Let (α_2, β_2) be the socially optimal outcome, and let (α_1, β_1) be the strategy in which each player betrays the other. In a Nash game, the only equilibrium is (α_1, β_1) , the worst possible outcome. In the limited-trust game, $\lim_{\delta \rightarrow 0} LTE(\delta)$ is (α_1, β_1) , but as δ increases, it shifts to (α_2, β_2) . [Table 5](#) shows the general form of a symmetric version of the game, with $0 < d_1 < c < d_2$. By noting the fact that both players will be playing the same strategy $\sigma = \{p_1, 1 - p_1\}$ at equilibrium if $\delta_1 = \delta_2$, we can find the $LTE(\delta)$ by solving the quadratic equation $(1 - p_1)(p_1 d_1 + (1 - p_1)d_2 - (1 - p_1)c) = \delta_1$ which yields

$$p_1 = \frac{2(d_2 - c) - d_1 \pm \sqrt{(d_1 - 2(d_2 - c))^2 - 4(d_2 - c - d_1)(d_2 - c - \delta_1)}}{2(d_2 - c - d_1)}$$

provided $d_2 - c - d_1 \neq 0$. If $d_2 - c - d_1 = 0$ then $p_1 = \frac{\delta_1 - d_2 - c}{d_1 - 2(d_2 - c)}$.

3. Computation of 2-Player $LTE(\delta)$

In this section we present a mathematical program for computation of an $LTE(\delta)$ in 2-player games. However, before doing so we consider the concept of a *dominated strategy* in a limited-trust game. By removing strongly dominated strategies, we will make the game smaller to aid in computation.

In a Nash game, a pure strategy s for player i is said to be dominated if $u_i(s, \sigma_{-i}) \leq u_i(s', \sigma_{-i}) \forall \sigma_{-i}$ for some alternate feasible strategy s' which is a convex combination of player i 's other pure strategies. s' is said to *weakly dominate* s if there is at least one σ_{-i} for which there is equality and at least on for which there is strict inequality. It is said to *strictly dominate* s if there is strict inequality for all σ_{-i} . In a limited-trust game with given δ , s' is said to dominate s if $u(s, \sigma_{-i}) \leq u(s', \sigma_{-i}) \forall \sigma_{-i}$ and $u_i(s, \sigma_{-i}) \leq u_i(s', \sigma_{-i})$ for all σ_{-i} . Weak dominance occurs if there is some σ_{-i} for which $u(s, \sigma_{-i}) = u(s', \sigma_{-i})$. As with Nash equilibria, no LTE will have a player i playing a strictly dominated pure strategy s_i with nonzero probability. Also as in Nash equilibria, we can iteratively remove dominated strategies by examining each strategy individually to see if it is dominated by a convex combination of the other still present strategies (this is done by using a linear program).

Having introduced the idea of dominance in the limited-trust context, we now demonstrate it on the game in [Table 6](#). From a Nash perspective, it is clear that α_2 strictly dominates α_3 and α_4 and weakly dominates α_1 . Similarly, β_1 strictly dominates β_3 . Therefore, these strategies need not be considered when looking for an MNE. From a limited-trust perspective this changes. α_2 no

Table 6
A game with Nash and limited-trust dominated strategies.

		Player 2		
		β_1	β_2	β_3
Player 1	α_1	0,7	5,5	0,5
	α_2	3,2	5,4	7,1
	α_3	0,6	4,1	1,5
	α_4	2,1	3,10	1,0

Table 7
Equivalent to the game in [Table 6](#) for MNE's (left) and $LTE(\delta)$'s (right).

		Player 2				Player 2	
		β_1	β_2			β_1	β_2
Player 1	α_1	0,7	5,5	Player 1	α_1	0,7	5,5
	α_2	3,2	5,4		α_2	3,2	5,4
					α_4	2,1	3,10




Fig. 4. Image in which s_1 is part of limited-trust best response and s_2 is not.

longer dominates α_1 , α_3 , or α_4 and β_1 no longer dominates β_3 . Although no pure strategy dominates α_3 , consider $\sigma_1 = [\frac{3}{5}, \frac{2}{5}, 0, 0]$. $u_1(\sigma_1, \beta_i) > u_1(\alpha_3, \beta_i) > u(\alpha_3, \beta_i)$ for $i \in \{1, 2, 3\}$ so α_3 is still strictly dominated and can be dropped from the problem. However, α_4 is part of the socially optimal σ and therefore cannot be strictly dominated unless β_2 is strictly dominated first. While β_3 cannot be strictly or weakly dominated by a convex combination of β_1 and β_2 in the original problem, consider the problem after α_3 is removed. For a mixed strategy $\sigma_2 = [\frac{1}{5}, \frac{4}{5}, 0]$ we see that $u_2(\alpha_i, \sigma_2) > u_2(\alpha_i, \beta_3)$ and $u(\alpha_i, \sigma_2) > u(\alpha_i, \beta_3)$ for $i \in \{1, 2, 4\}$. Therefore, while we cannot remove β_3 immediately as in the Nash case, we can still remove it through the iterated removal of other dominated strategies. In the Nash case we then get the equivalent game in the left side of [Table 7](#), and for the limited-trust case we get the equivalent game on the right side.

It is interesting to note that the value of δ is not relevant in determining whether a strategy is dominated in a game. This is because we cannot say that s' dominates s if $u(s, \sigma_{-i}) \leq u(s', \sigma_{-i}) \forall \sigma_{-i}$ and either $u_i(s, \sigma_{-i}) \leq u_i(s', \sigma_{-i})$ or $u_i(\sigma_i^C, \sigma_{-i}) - u_i(s', \sigma_{-i}) \leq \delta_i$ for all σ_{-i} . If the second condition occurs, s may still be part of a unique limited-trust best response. An example of this is given in [Fig. 4](#), in which, for a fixed σ_{-i} , player i has three strategies s_1, s_2 , and s_3 : despite the fact that $u(s_2, \sigma_{-i}) > u(s_1, \sigma_{-i})$

and $u_i(\sigma_i^G, \sigma_{-i}) - u_i(s_2, \sigma_{-i}) < \delta_i$, the limited-trust best response is a convex combination of s_1 and s_3 , but not s_2 .

Having defined limited-trust dominance to reduce computational effort, we now introduce our solution method. Our mathematical program for finding an LTE in a 2-player bimatrix game given by $A, B \in \mathcal{R}^{m \times n}$ will be loosely based on the linear program used in the support enumeration algorithm for finding Nash equilibria, which determines if a given support pair S_A, S_B admits an MNE (i.e. there is an MNE in which only the pure strategies in S_A, S_B are played with positive probability, and all such strategies are played with positive probability). The mathematical program is given by MP1 where

$$f_x(y) = \max_x x^T (A + B)y \quad \text{subject to} \quad x^T A y \geq e_j^T A y - \delta_1, \quad \forall j \in [m]$$

and $f_y(x)$ is similarly defined. MP1 constitutes a quadratically-constrained program with a bilevel component from $f_x(y)$ and $f_y(x)$. A, B are the $m \times n$ payoff matrices for players 1 and 2, respectively, and e_j is the vector with the value 1 at index j and zero elsewhere.

MP1:

$$\begin{aligned} & \max_{x,y} \quad 1 \\ & \text{subject to} \\ & \quad x^T A y \geq e_j^T A y - \delta_1 \quad \forall j \in [m] \quad (1) \\ & \quad x^T B y \geq x^T B e_j - \delta_2 \quad \forall j \in [n] \quad (2) \\ & \quad x^T (A + B)y \geq f_x(y) \quad (3) \\ & \quad x^T (A + B)y \geq f_y(x) \quad (4) \\ & \quad \sum_{i \in S_A} x_i = 1 \quad (5) \\ & \quad \sum_{j \in S_B} y_j = 1 \quad (6) \\ & \quad x_i = 0 \quad \forall i \notin S_A \quad (7) \\ & \quad y_j = 0 \quad \forall j \notin S_B \quad (8) \\ & \quad x_i \geq 0 \quad \forall i \in [m] \quad (9) \\ & \quad y_j \geq 0 \quad \forall j \in [n] \quad (10) \end{aligned}$$

The bilevel elements of constraints (3) and (4) are a necessary portion of the program: without these constraints, which force the total utility to be the greatest possible when each player is playing within δ_i of its greedy best response to the other, the solutions of MP1 would simply be a subset of the ε -equilibria for $\varepsilon = \max_i \delta_i$, regardless of whether they were also limited-trust equilibria. This cannot be solved using the objective function to drive the program, as the socially optimal ε -equilibrium is not necessarily an LTE. We also mentioned above that MP1 is loosely based on the Support Enumeration algorithm for finding 2-player MNEs. However, due to the fact that a general $\text{LTE}(\delta)$ is not a well-supported equilibrium, we are unable to fully linearize the constraints as in the support enumeration algorithm for finding MNE's. As a consequence, if $S_A \subseteq S_C$ and $S_B \subseteq S_D$, then any solution to $\text{MP1}(S_A, S_B)$ is also a solution to $\text{MP1}(S_C, S_D)$, which is not the case in the Nash support enumeration. The same problem is observed in finding non-well-supported approximate Nash equilibria as well, so this is not surprising.

We now prove below any $\text{LTE}(\delta)$ given by (x, y) is a solution to MP1 for appropriate S_A, S_B .

Theorem 7. A strategy set (x, y) for a two player game is an $\text{LTE}(\delta)$ if and only if it is a feasible solution to MP1 for $S_A = [m], S_B = [n]$.

The proof of this Theorem may be found in [Appendix A.2](#).

Corollary 8. For any feasible solution to MP1, constraints (3) and (4) are fulfilled with equality.

Proof. Follows from [Theorem 7](#): any solution to MP1 is an $\text{LTE}(\delta)$, and any $\text{LTE}(\delta)$ fulfills the constraints with equality as each player is playing a limited-trust best response to the other. $\square$

Given that MP1 was stated to have been loosely based on the linear program used in the support enumeration algorithm for Nash equilibria, it is natural to question why the program is not set up to iterate over supports, as in that algorithm. This comes about because in any greedy best response, every pure strategy which player i plays against the other player is a best response, and so the quadratic constraints (1) and (2) in MP1 can be transformed into a larger set of linear constraints which enforce the condition that every pure strategy in the support of a Nash equilibrium is a greedy best response. There is no corresponding condition for an LTE which allows us to consider the pure strategies of a support individually rather than the mixed strategy LTEs as a whole. However, if we are looking for well-supported LTE's we can use a support enumeration method by replacing constraints (1) and (2) in MP1 with those below and then apply [Algorithm 1](#).

$$\begin{aligned} & e_i^T A y \geq e_i^T A y - \delta_1 \quad \forall i \in S_A \quad \forall j \in [m] \\ & x^T B e_i \geq x^T B e_j - \delta_2 \quad \forall i \in S_B \quad \forall j \in [n] \\ & x_i = 0 \quad \forall i \notin S_A \\ & y_j = 0 \quad \forall j \notin S_B. \end{aligned}$$

[Algorithm 1](#) will find at least one LTE for every support pair S_A, S_B which admits a well-supported LTE. However, as we have already seen well-supported LTE's may not exist.

Algorithm 1 $\text{LTESupportEnumeration}(A, B, \delta_1, \delta_2)$.

```
Initialize hashset LTESet  $\leftarrow \emptyset$ ;
for  $S_A \in [m], S_B \in [n]$  do
   $(x, y) \leftarrow \text{SolveMP1}(S_A, S_B)$ ;
  if  $(x, y) \notin \text{LTESet}$  then
     $\text{LTESet}[S_A, S_B] \leftarrow (x, y)$ 
return LTESet
```

4. Leader-Follower Equilibria

We have defined the concept of Limited-Trust equilibria in simultaneous games in a natural manner, and showed that at least one LTE exists in any simultaneous game of n players. The next natural extension to consider is LTE's in turn-based games, i.e. leader-follower or *Stackelberg* games.

Consider a two-player turn-based game of complete information, i.e. player 1 picks from m strategies and in response, player 2 picks from n strategies with full knowledge of the first player's choice. Such a game is akin to a bi-level optimization problem for the first player: given full-knowledge by all players, the second player's response is deterministically dictated by the first player. As such, this game always has a pure equilibrium known as the Stackelberg equilibrium and, assuming a fixed tie-breaking rule for players between multiple equivalent strategies, the Stackelberg equilibrium is unique. While this is true for n -player games, for the sake of simplicity we will confine our discussion to $n = 2$, as $n > 2$ follows naturally.

We now want to consider what happens when players have trust levels δ_1, δ_2 . Although this is a full knowledge deterministic game with regard to the payoffs and the first player's action being known to the second player, unlike in the simultaneous game the nature of the equilibrium changes sharply depending on the first player's knowledge of δ_2 and the values from which each player measures δ_i . Because of this, we will examine three policies which represent different interpretations of a Limited-Trust Stackelberg

Equilibrium (LTSE). We assume a bimatrix game with payoff matrices $A, B \in \mathcal{R}^{m \times n}$ for players 1 and 2, respectively. Note that here n is the number of pure strategies possessed by the second player.

1. Incomplete Knowledge: The first player does not know anything about δ_2 and, being risk averse, assumes the second player is not trustworthy ($\delta_2 = 0$). The first player then determines the second player's response to each of its possible actions under this assumption and finds strategy i such that $i = \arg \max_{0 \leq j \leq m} a_{jr(j)}$ where $r(j)$ is player 2's best response to j and a_{ij} and b_{ij} are the first and second players' payoff if they play i and j , respectively. The first player then plays j which maximizes $a_{jr(j)} + b_{jr(j)}$ subject to $a_{ir(i)} - a_{jr(j)} \leq \delta_1$, and the second player plays l which maximizes $a_{jl} + b_{jl}$ subject to $b_{jr(j)} - b_{jl} \leq \delta_2$.
2. Complete Knowledge: The first player knows δ_2 . It knows that if it plays i , then player 2 will play its best response $s(i)$ which maximizes $a_{is(i)} + b_{is(i)}$ subject to $b_{ir(i)} - b_{is(i)} \leq \delta_2$. Player 1 then finds i such that $i = \arg \max_{0 \leq j \leq m_1} a_{js(j)}$, and plays j which maximizes $a_{js(j)} + b_{js(j)}$ subject to $a_{is(i)} - a_{js(j)} \leq \delta_1$. The second player then plays $s(j)$.
3. Cooperative Complete Knowledge: Let i, j be the regular Stackelberg Equilibrium. The players play k, l which maximizes $a_{kl} + b_{kl}$ subject to $a_{kl} - a_{ij} \leq \delta_1$ and $b_{kl} - b_{ij} \leq \delta_2$.

Of these policies, the first two seem like the most natural interpretations of the LTE in the turn-based game: a player is willing to forgo a payoff at most δ_i higher than what they could get, provided that the other player gains at least that much. The only question is whether or not player 1 knows δ_2 : while the question was unimportant in the simultaneous setting as equilibrium was merely a point where no player could unilaterally improve the total utility without exceeding its maximum cost, here the leader-follower nature of the game means the first player can determine exactly how the second player will act and plan its strategy accordingly. The only question for the first player is the value of δ_2 , as if it's not sure then it must plan for the worst and assume $\delta_2 = 0$.

While the cooperative complete knowledge policy may seem less natural, the confusion is a matter of perspective: with the first player having full knowledge of δ_2 , instead of measuring its payoffs over the second player's reactions to each strategy i it could play, it instead measures them with respect to the Stackelberg equilibrium. The second player makes the same choice: it is rational and can determine the first player could have played according to the Stackelberg equilibrium if it wished to, and so reacts accordingly. This policy requires more coordination between players, but can be interpreted as two players who regularly interact and strive to maintain a good relationship. In this sense it is less suited for one-off games. However, the same could be said of the complete knowledge policy, as it is otherwise infeasible to expect the first player to know δ_2 *a priori*.

We have derived additional results in regard to the expectations and probabilities for all three of these policies in random leader-follower games. However, the details are somewhat involved and do not provide any great insight to the reader. As such, these results and their derivations are available in [Appendix B](#).

5. Numerical results

In this section we present a numerical comparison of the efficiencies of the LTE when compared to Nash equilibria in both the simultaneous and leader-follower settings. LTE's are found for randomly-generated games and compared with the maximum-value Nash and Stackelberg PoA's of these games. These represent random repeated games, a set of games which, while not identical, are all drawn from the same distribution. These games are of particular interest because while the LTE is explicitly created for ana-

lyzing one-off games, it is implicitly motivated by the expectation that future games will be played. Day-to-day societal interactions are a perfect example of this, and are well modeled by random repeated games: such interactions between players are not identical, but will display a pattern over time so that they could be said to come from some "typical" distribution.

Theoretical results related to random repeated games in the leader-follower setting can be viewed in detail in [B.3](#), but one which we will state here is that the expected PoA of the Stackelberg equilibrium for a 2-player leader-follower bimatrix game. In such a game with $m \times n$ payoff matrices where each entry generated is generated independently and from an identical distribution (iid) A or B for players 1 and 2, respectively, the expected PoA is

$$\frac{E[(A+B)^{(mn)}]}{E[A^{(m)}] + E[B^{(n)}]}$$

where $X^{(n)}$ is the maximum of n samples x_i generated iid from distribution X .





We consider random repeated games in our numerical trials. These are represented as bimatrix games where the payoff matrices A, B for each player are generated according to some distribution. In particular, we consider matrices where each entry is generated iid for each player, though A and B may not come from the same distribution. The majority of games were generated as 2-player 2×2 repeated games, with entries of player 1's payoff matrix generated iid according to a distribution A and player 2's payoff matrix generated iid according to a distribution B . One instance of 2-player 3×3 games was generated, as unlike in the leader-follower games we have a less precise bound on the PoA as a function of m and n . LTE's of simultaneous games were computed using MP1 in [Section 3](#), and LTSE's were computed for each leader-follower policy using the methods given in their descriptions in [Section 4](#).

5.1. Simultaneous Games

We consider games where players' payoffs are generated iid from three distributions: geometric with $p = \frac{1}{4}$, uniform over the set of integers in $[0,10]$, and Normal $\mathcal{N}(0, 1)$. 100 instances of 2×2 games are generated from each of these distributions, as well as 100 3×3 games drawn from the geometric distribution. We then vary $\delta_1 = \delta_2 = \delta$ from 0.01 to 1 for each game, to see how the value of the LTE changes as a function of δ . For each test case, we use the support enumeration method to determine the MNE with the highest net utility for comparison to the LTE with the highest net utility. Because any MNE which provides the optimal level of social utility is also an LTE, we ignored generated test cases where the social optimum was also an MNE.

Because of the nonconvex and potentially disjoint nature of the solution space, we include a constraint in MP1 that the value of the net utility of the LTE must be greater than or equal to the utility provided by the best MNE (net utility-maximizing MNE). As we have already seen in [Section 2](#), such an LTE may not exist for all values of δ . Therefore, if the solver fails to locate a feasible solution to MP1 after 50 attempts on a particular test case, this constraint is relaxed. It is not reintroduced until after δ has increased to a level where an LTE with better net utility than the MNE is discovered. Additionally, while we compute an LTE which has higher net utility than all MNE's, it may not necessarily be the maximum value LTE. This is due to the nonconvexity of the set of LTE's.

[Figs. 5 through 8](#) detail the results of our numerical trials. Although δ is always varied from 0.01 to 1, in the figures it is rewritten as the percentage of the net utility generated by socially optimal (net utility-maximizing) MNE so as to compare values across different distributions. In all but [Fig. 7](#), we see a very clear linear

Fig. 5. 2×2 Geometric Games.Fig. 7. 2×2 Normal Games.Fig. 6. 3×3 Geometric Games.Fig. 8. 2×2 Uniform Games.

relationship between δ as a percentage of the maximum-valued MNE and LTE as a percentage of the MNE. In both the geometric games, the curve has a slope of approximately 2, meaning that on average, for every game a player has to give up δ , there is a game where it gains 3δ over what it would receive by playing selfishly. The uniform games in Fig. 8 show a similar result, with a slope of approximately 1.3.

This brings us to Fig. 7, which unlike the others does not evince an approximately linear curve. However, consider the variance of the distributions: the geometric distribution with $p = \frac{1}{4}$ has a variance of 12 and the discrete uniform distribution over $[0,10]$ has a variance of 10. In contrast, the variance of 1 possessed by the Normal distribution is quite small. Now consider what the curves in Figs. 5, 6, and 8 would look like if we continued to increase δ : the curves would eventually start to evince diminishing returns, as increasing δ past the point where many social optima start to become LTE's will produce very little additional social utility. This explains the curve in Fig. 7: it is merely a curve in which the δ is already quite large compared to the variance of the distributions from which the entries of A and B are generated, and thus is experiencing diminishing returns. It also indicates that if we continue to increase δ , the other figures will come to resemble Fig. 7.

5.2. Leader-Follower Games

We conduct numerical studies on 2×2 games with payoff matrices generated from three distributions: $U[-0.5, 0.5]$, $\mathcal{N}(0, 1)$, and $\exp(1)$. For each set of trials, we let $A \sim B$, and let $\delta_1 = \delta_2 = \delta$. We define the Stackelberg gap as the difference between the Stackelberg equilibrium and the social optimum. Figs. 9, 10, and 11 each show the average PoA of 1000 games generated according to these distributions and solved for $\delta \in [0, 1]$ in the first graphs, where as mentioned in Section 4, $\delta = 0$ indicates that there is no trust between the players. The second graph considers how much of the Stackelberg gap is covered by each of the policies at the varying δ levels. Unlike in the simultaneous case, in the leader-follower setting under the complete knowledge policy if the Stackelberg equilibrium is the social optimum that does not guarantee it is also the LTSE. For that reason we have not ignored games in which this occurs.

Unsurprisingly, in all three distributions for all values of δ , the cooperative complete knowledge policy results in the best performance. We even note that with $\delta = 0$ it still manages to recover an average of approximately 20% or greater of the Stackelberg gap for each tested set of games. This "Reward without Risk" comes from the greater cooperation between players seen in this interpretation of the LTE.




Fig. 9. Leader-Follower Numerical Results, $A \sim B \sim U[-0.5, 0.5]$.




Fig. 10. Leader-Follower Numerical Results, $A \sim B \sim \mathcal{N}(0, 1)$.

Also unsurprisingly, the complete knowledge policy tends to outperform the incomplete knowledge policy on average, for most δ values. Fig. 9 provides an excellent demonstration of this: because the max and min possible payoffs have a gap of 1, by the time $\delta = 1$ both players are trying only to maximize the social utility. In particular, by the time δ reaches approximately 0.75, the complete knowledge game tends to result in the social optimum

being played virtually every time. This is because the entries of A and B are drawn from $U[-0.5, 0.5]$ so the chance of the socially optimal outcome having a utility for player 1 which is more than 0.75 less than the player's greediest move is nearly 0. In contrast, while the same is true in the incomplete knowledge case, the game levels off to covering slightly over 80% of the Stackelberg gap even at $\delta = 1$. This occurs due to the fact that although both players are




Fig. 11. Leader-Follower Numerical Results, $A \sim B \sim e^{-x}$.

effectively altruistic at this level of δ , the first player does not believe that the second player is. This causes player 1 to attempt to maximize the social utility around the assumption that $\delta_2 = 0$, despite the fact this is not true. We can consider this gap between the incomplete and complete knowledge cases as the cost of ignorance.

It is important to note that the cost of ignorance may not be bad. Indeed Fig. 11 shows that for δ between approximately 0.05 and 0.3, the cost of ignorance is negative. This occurs due to the fact that the first player is unaware that the second player is willing to give up δ , and thus is unable to take advantage of that fact for its own gain. This is identical to what happens for some values of δ in the game described in Table B.8 in Appendix B.1.

6. Discussion and summary

Throughout this paper we have been considering limited-trust equilibria as a description of behavior which is not entirely selfish, provided the opportunity cost of the behavior for player i is less than some bound δ_i . This idea of an LTE was expressed very naturally in simultaneous games, where at equilibrium each player does not care about the δ values which are motivating other players, only that it plays its best response to what those players are actually doing. The key managerial insight of the LTE is that while the players in giving something up appear to be playing “non-rationally” when games are considered in isolation, when considered as a whole both players actually achieve more than they would have received if they had myopically played the “rational” Nash equilibria in each game. We saw that while it was possible for limited-trust games to have worse results than Nash games, it will not happen in several common classes of games, and occurs rarely in others: in 2-player numerical trials with $\delta_1 = \delta_2 = \delta$ we observed an average personal utility increase of δ for each player when δ was modest compared to the variance in the utilities of randomly generated games. When we consider the leader-follower setting players can no longer ignore the δ values of their fellows,

and we considered the effects of whether or not players knew each other's δ 's or had to prepare for the worst (assume $\delta_{-i} = 0$).

It is natural to question the method developed in this paper for the computation of LTE's in a simultaneous game. Given that such LTE's are a subset of ε -equilibria, which are PPAD-hard to compute for general ε , we do not expect to derive an algorithm for the general k -player case without going to a mathematical program similar to MP1. However, readers may wonder why we have not provided a different algorithm for computing an LTE in the 2-player game.

The Lemke-Howson algorithm (Lemke & Howson (1964)) is one of the first algorithms for finding Nash equilibria in a 2-player bimatrix game and remains one of the most popular. It relies on the observation that at equilibrium (σ_1^G, σ_2^G) , if player i has m_i pure strategies then for a best response $\sigma_i^G = \{p_1^i, p_2^i, \dots, p_{m_i}^i\}$ either $p_j^i = 0$ or playing the pure strategy s_j^i is a best response to σ_{-i} . With this observation, the Lemke-Howson algorithm is able to set up a linear complementarity program (LCP) for which any feasible solution is a Nash equilibrium. Unfortunately the definition of an LTE does not lend itself well to this method. This is partially due to the fact that we cannot make a similar observation about pure strategies in an LTE. However, while this problem may be possible to overcome, the larger difficulty comes from the fact that there is an optimization problem embedded in each player's best response to the other. While this is also true of a greedy best response, that optimization problem can be expressed solely as a set of linear constraints with no objective function, i.e. σ_i is a best response to σ_{-i} if and only if $u_i(\sigma_i, \sigma_{-i}) \geq u_i(e_j, \sigma_{-i})$ for all $j \in [m_i]$. The optimization problem embedded in the limited-trust best-response cannot be absorbed to a larger program due to having an objective function. This is reflected in the fact that the best response function is explicitly brought into MP1 in constraints (3) and (4), rather than bringing in constraint sets. Even the further generalization of the algorithm in Lemke (1965) is unlikely to adapt to computing LTEs. Although the Lemke-Howson algorithm is nearly

sixty years old and has since been shown to be a special case of the Global Newton Method by Govindan and Wilson (2003), it remains an extremely prevalent method for computing Nash equilibria in 2-player finite games in practice. This is particularly true following the proof by Chen and Deng (2006) that ε -equilibria (and Nash equilibria) are PPAD-complete to compute even for 2-player finite games. It is worth noting that as a consequence, $\text{LTE}(\delta)$ is also PPAD-hard to compute.

We also considered several natural interpretations of the LTE in a leader-follower game, which vary drastically depending on how much knowledge players have of each other. More definite theoretical probabilities for the likelihood of a social optimum occurring in a random game in the leader-follower context, as these equilibria are significantly easier to compute. We then moved onto numerical testing of the LTE, comparing how the social utility varied over random repeated games as a function of δ , particularly when compared to Nash and Stackelberg equilibria. One of the more surprising results of our numerical trials in simultaneous games was how strong the linear relationship was between the net utility and δ , prior to the onset of diminishing returns as δ continues to increase. In our leader-follower games we observed the differences in the utility of each of our interpretations, noting that the cooperative complete knowledge case produced significant gains at the no risk level of $\delta = 0$, and also that the gap between complete and incomplete knowledge effectively measured the price of ignorance. Perhaps more surprising was that the price of ignorance was sometimes negative on average, rather than just occasionally, with parameters existing for which player 1 assuming the worst of player 2 resulted in higher average total utility.

As noted earlier in this section, while many traditional equilibrium computation methods such as the Lemke-Howson algorithm are unlikely to adapt well to the LTE computation, we would still like to put more study into the computation of simultaneous game LTE's. Additionally, we are interested in considering how LTE's model behavior in larger systems such as social networks. Perhaps the most exciting line of inquiry is that of learning: the LTE is positioned as a tool for non-Nash analysis of repeated game that can also solve one-off simultaneous games, something for which there are few existing tools. As such each player should be trying to set their δ_i in order to maximize their utility over time. We are very interested in the potential dynamics of shifts in δ values as players interact with each other, particularly if they take each other's playing history into account. Also of interest is the relationship between talented but relatively selfish individuals as opposed to trustworthy individuals without specialized skills, and the resulting "diva" behavior often exhibited by the former. We will focus on these areas of study in our future research.

Acknowledgement

The work of J. Garg was supported by NSF CRII under Award 1755619. T. Murray and R. Nagi were supported in part by ONR through the Program Management of Drs. D. Wagner and W. Adams under Award N000014-16-1-2245.

Appendix A. Proofs

A.1. Proof of Lemma 1

Proof. Consider that

$$\sigma_i^* = \arg \max_{\sigma_i \in \Sigma_i} u(\sigma_i, \sigma_{-i}^*)$$

subject to

$$\delta_i \geq u_i(\sigma_i^G, \sigma_{-i}^*) - u_i(\sigma_i, \sigma_{-i}^*)$$

prior to the constant c_j being added. Each player $i \neq j$ now has to solve

$$\sigma_i^{*'} = \arg \max_{\sigma_i \in \Sigma_i} (u(\sigma_i, \sigma_{-i}^*) + c_j)$$

subject to

$$\delta_i \geq u_i(\sigma_i^G, \sigma_{-i}^*) - u_i(\sigma_i, \sigma_{-i}^*)$$

which has the same solution, and player j has to solve

$$\sigma_j^{*'} = \arg \max_{\sigma_j \in \Sigma_j} (u(\sigma_j, \sigma_{-j}^*) + c_j)$$

subject to

$$\delta_i \geq (u_j(\sigma_j^G, \sigma_{-j}^*) + c_j) - (u_j(\sigma_j, \sigma_{-j}^*) + c_j)$$

which also has the same solution. Therefore, σ^* is still an $\text{LTE}(\delta)$. $\square$

A.2. Proof of Theorem 7

Proof. First, suppose that we have a feasible solution (x, y) . Given that it is a feasible solution, the last four constraints will not be violated as x and y are valid strategy profiles. The first two constraints ensure that neither player is giving up more than δ_i that it could be making by playing the greedy best response to the other player's strategy. The next two ensure that the social utility from the players actions is at least the amount which would be provided if each player played its limited trust best response to the other. Therefore, since each player is providing as much social utility as if it had been playing its limited trust best response (more is impossible without violating constraint 1 or 2) and is not giving up more than δ_i of what it could be making, (x, y) is an $\text{LTE}(\delta)$.

Now, consider an $\text{LTE}(\delta)$ given by (x, y) , which we will show to be a feasible solution. Because (x, y) is an LTE, x and y are both valid strategy profiles and thus do not violate constraints (5-10), particularly as constraints (7) and (8) are disabled for $S_A = [m]$, $S_B = [n]$. (x, y) is an LTE, so neither player is giving up more than δ_i and therefore constraints (1) and (2) are fulfilled. Finally, because it is an LTE, x is a limited-trust best response to y and y is a limited-trust best response to x , which means that constraints (3) and (4) are fulfilled with equality. $\square$

Appendix B. Additional Results: Leader-Follower Games

B.1. Demonstration of Policies

Given the interpretation of the policies in Section 4, we now provide a practical demonstration of each on the 2×2 game given in Table B.8. The Stackelberg equilibrium in this game occurs when the first player plays 2 and the second player plays 1. We first consider the incomplete knowledge policy: for $\delta_1 < 1$, the first player will play 2, as it otherwise stands to lose $5 - 4 = 1$ if it plays 1. If the first player plays 1, then for $\delta_2 < 1$, the second player plays 2, but for $\delta_2 \geq 1$, the second player plays 1 as well. If the first player plays 2, then the second player plays 1 for $\delta_2 < 1.5$ and 2 for $\delta_2 \geq 1.5$. Thus for the socially optimal policy (1,1) to be played, we must have $\delta_1, \delta_2 \geq 1$.

Table B.8
A 2-Player 2×2 Game.

		p_2	
		1	2
p_1	1	6, 4	4, 5
	2	5, 3	8, 1.5

We next consider the complete knowledge policy. The second player's behavior is the same: if the first player plays 1 the second plays 1 for $\delta_2 \geq 1$ and if the first player plays 2 the second plays 2 for $\delta_2 \geq 1.5$. Suppose $\delta_2 < 1$, then the first player is selecting between (2,1) and (1,2): if $\delta_1 \geq 1$ then it plays 1, and otherwise it plays 2. Suppose $1 \leq \delta_2 < 1.5$, then the first player is choosing between (1,1) and (2,1) and so plays 1 regardless of δ_1 . Finally, suppose $\delta_2 \geq 1.5$. Here, the first player is selecting between (1,1) and (2,2). If $\delta_1 < 2$, the first player plays 2 to receive 8 after the second player also plays 2. If $\delta_1 \geq 2$ it plays 1 instead.

The complete knowledge policy has the property that if we have $(\delta'_1, \delta'_2) \geq (\delta_1, \delta_2)$, that does not mean that the net utility for the (δ'_1, δ'_2) game is greater than or equal to that of the (δ_1, δ_2) game: the $\delta_1 = \delta_2 = 1$ game results in a greater net utility than the $\delta_1 = \delta_2 = 1.5$ game over the utility above. This is interpreted as the first player taking advantage of the second player's trustworthiness. If $(\delta'_1, \delta'_2) \geq (\delta_1, \delta_2)$ then we can say that at least one of the following is true: player 1 receives higher utility under (δ'_1, δ'_2) , or the net utility is greater under (δ'_1, δ'_2) . This is because had player 1 selected the same strategy as it would with (δ_1, δ_2) , it would have achieved at least as much utility for $\delta'_2 \geq \delta_2$. The fact that it did not select the same strategy indicates that it selected one which increases either its own utility or the net utility.

Finally, we consider the cooperative complete knowledge policy. This policy is especially interesting as for any $\delta_1, \delta_2 \geq 0$, the social optimum of (1,1) will be played: as stated earlier, the Stackelberg equilibrium in this game occurs when the first player plays 2 and the second player plays 1. The social optimum occurs at (1,1), with a net payoff of 10, and at the social optimum both players are receiving strictly more than they would in the Stackelberg equilibrium. This is the only policy which has the possibility of reward without risk: for any game in which $\delta_1 = \delta_2 = 0$, each player is guaranteed a minimum of what they would achieve in the Stackelberg equilibrium and the possibility of more. For the other two policies, $\delta_1 = \delta_2 = 0$ guarantees that they will play the Stackelberg equilibrium, meaning no risk, no reward.

Because the behavior of games under the cooperative complete knowledge policy is different from traditional leader-follower games even when $\delta_1 = \delta_2 = 0$, we will use the term *zero-trust* game to refer to the traditional leader-follower game.

B.2. Playing Social Optima

In this subsection we will examine what the structure of a 2-player bimatrix game must be in order for the Socially Optimal (greatest net payoff) strategy combination to be played. We will confine our discussion to 2×2 games, but the result generalizes to $m \times n$ games.

Let A, B be the payoff matrices for the first and second players, respectively. Without loss of generality, assume that the social optimum occurs when (1,1) is played.

First, we consider what must happen for the Stackelberg equilibrium to be the social optimum. An immediate requirement is $b_{12} \leq b_{11}$, as otherwise even if the first player plays 1, the second player will play 2. Without loss of generality, assume $b_{21} \geq b_{22}$ so that if the first player plays 2, the second player will play 1. In order for the first player to play 1 rather than 2 we must have $a_{11} \geq a_{21}$. More formally, (1,1) is the social optimum and gets played if and only if E_1, E_2, E_3, E_4 are satisfied where

$$E_1 = (a_{11} + b_{11} \geq \{a_{12} + b_{12}, a_{21} + b_{21}, a_{22} + b_{22}\})$$

$$E_2 = (b_{2i} \geq \{b_{21}, b_{22}\})$$

$$E_3 = (a_{2i} \leq a_{11})$$

$$E_4 = (b_{11} \geq b_{12}).$$

and i is player 2's best response to the first player playing 2.




Fig. B.12. Geometric Leader-Follower Representation.

Geometrically, we can determine the outcome of a game by plotting its entries in $\mathcal{R}^2$ in terms of the utility for each player, as in Fig. B.12 where the horizontal axis a is the utility for the first player and the vertical axis b is the utility for the second player.

From the figure, we can see that (1,1) is the social optimum. It is also the Stackelberg equilibrium: $b_{12} < b_{11}$, so the second player will play 1 if the first plays 1. $b_{21} > b_{22}$, so the second player will play 1 if the first plays 2, and $a_{21} < a_{11}$, so the first player will play 1. In terms of the geometry, we can say that (1,1) is the Stackelberg equilibrium and the social optimum if and only if (1,2) is not in area 1 and whichever is larger out of b_{22} or b_{21} , that point is not in area 3. Then the first player will prefer 1 to 2 and the second will prefer (1,1) to (1,2). This is true for all complete knowledge 2×2 leader-follower games, and is easily generalized to $m \times n$ games.

Now consider the incomplete knowledge policy. If the social optimum occurs at (1,1), what must occur for it to be played? If $b_{11} \geq b_{12}$ and $b_{2i} \geq b_{21}, b_{22}$ the first player will play 1 if $a_{2i} \leq a_{11} + \delta_1$, and the second player will also play 1. If $b_{11} < b_{12}$ and $b_{2i} \geq b_{21}, b_{22}$, then the first player will play 1 if $a_{2i} < a_{12} - \delta_1$ or if $a_{2i} \leq a_{12} + \delta_1$ and $a_{2i} + b_{2i} < a_{12} + b_{12}$ or $a_{2i} \geq a_{12} + \delta_1$. If the first player plays 1, the second player will play 1 as well if $b_{12} \leq b_{11} + \delta_2$. We can therefore say that (1,1) is the social optimum and gets played if and only if $E_1, E_2, F_3, E_4, E_1, E_2, F_5, F_6, F_7$, or $E_1, E_2, F_5, -F_6, F_8$ are satisfied where

$$F_3 = (a_{2i} \leq a_{11} + \delta_1)$$

$$F_5 = (b_{11} < b_{12} \leq b_{11} + \delta_2)$$

$$F_6 = (a_{2i} + b_{2i} < a_{12} + b_{12})$$

$$F_7 = (a_{2i} \leq a_{12} + \delta_1)$$

$$F_8 = (a_{2i} \leq a_{12} - \delta_1)$$

As with the Stackelberg game, we geometrically model these constraints by plotting the payoffs in $\mathcal{R}^2$, with both possible constraint sets seen in Figs. B.13 and B.14. On both figures, given the points (1,1) and (1,2), the social optimum is played if whichever is greater of b_{21}, b_{22} is in the shaded blue area.

We next consider the complete knowledge policy. Again we assume, without loss of generality, that (1,1) is the social optimum. Let $b_{2i} \geq b_{21}, b_{22} \geq b_{2j}$. We want to consider when the second player would play j given that the first player plays 2. This only occurs if $a_{2i} + b_{2i} < a_{2j} + b_{2j}$ and $b_{2i} < b_{2j} + \delta_2$. Now we can determine that the first player will only choose 2 (given the second player would play the social optimum if the first chose 1) if $a_{2i} > a_{11} + \delta_1$ or $a_{2j} > a_{11} + \delta_1$ and $a_{2i} + b_{2i} < a_{2j} + b_{2j}$ and $b_{2i} \leq b_{2j} + \delta_2$. Therefore, the social optimum gets played if and only if E_1, E_2, F_3, F_5




Fig. B.13. E_1, E_2, F_3, E_4 .Fig. B.14. E_1, E_2, F_5, F_6, F_7 .

Fig. B.15. Geometric Complete Knowledge Leader-Follower Representation.




Fig. B.16. Geometric Cooperative Complete Knowledge Leader-Follower Representation.

AND (G_9 OR G_{10} OR G_{11}) is satisfied where

$$G_9 = (a_{2j} \leq a_{11} + \delta_1)$$

$$G_{10} = (a_{2i} + b_{2i} \geq a_{2j} + b_{2j})$$

$$G_{11} = (b_{2i} \geq b_{2j} + \delta_2).$$

Suppose $b_{21} \geq b_{22}$. If the geometric representation of the incomplete knowledge policy primarily depended on (2,1) and was dictated by (1,2), here it is dependent on (2,2) and dictated by (2,1). Fig. B.15 displays this: given the (1,1), (1,2), and (2,1) the red region represents where the point (2,2) cannot be in order for the social optimum to be played. Additionally, we must have $b_{12} \leq b_{11} + \delta_2$ and $a_{21} \leq a_{11} + \delta_1$.

Finally, we consider the cooperative complete knowledge policy. It is easy to write the requirements for the social optimum to be played in terms of the Stackelberg equilibrium: if (1,1) is the social optimum and (i, j) is the Stackelberg equilibrium, (1,1) is played if and only if $a_{ij} \leq a_{11} + \delta_1$ and $b_{ij} \leq b_{11} + \delta_2$. Fig. B.16 demonstrates this: if (1,1) is the social optimum, it is played if and only if the Stackelberg equilibrium occurs in the red shaded area.

B.3. Leader-Follower Random Repeated Games

While the LTE is capable of analyzing one-off games, it is also a suitable tool for repeated games with a particular emphasis placed on day-to-day societal interactions. Such interactions between players are not identical, but will likely display a pattern over time so that they could be said to come from some “typi-

cal” distribution. Because of this, we now consider games where the payoff matrices A, B for each player are generated according to some distribution. In particular, we consider matrices where each entry is generated iid for each player, though A and B may not come from the same distribution. Let f_a and f_b be the probability distribution functions of A and B , respectively.

Theorem 9. In a 2×2 leader-follower game with zero trust where all of player 1's payoffs are a_{ij} iid and are independent from player 2's payoffs b_{ij} which are iid, the probability that Stackelberg equilibrium of the game is the social optimum is

$$P(S.O.) = 8 \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} M(a_{11}, b_{11}) f_a(a_{11}) da_{11} f_b(b_{11}) db_{11}$$

where

$$M(a_{11}, b_{11}) = \left(\int_{-\infty}^{a_{11}} \left(\int_{-\infty}^{a_{11}+b_{11}-a_{21}} G(a_{11}, b_{11}, b_{21}) f_b(b_{21}) db_{21} \right) f_a(a_{21}) da_{21} \right) \\ \cdot \left(\int_{-\infty}^{b_{11}} \left(\int_{-\infty}^{a_{11}+b_{11}-b_{12}} f_a(a_{12}) da_{12} \right) f_b(b_{12}) db_{12} \right) \\ G(a_{11}, b_{11}, b_{21}) = \left(\int_{-\infty}^{b_{21}} \left(\int_{-\infty}^{a_{11}+b_{11}-b_{22}} f_a(a_{22}) da_{22} \right) f_b(b_{22}) db_{22} \right)$$

Proof. Given that every entry in the payoff matrix is equally likely to be the social optimum, and equally likely to be played if it is the social optimum, we will approach this problem by finding the probability that 1,1 is the social optimum, and is played. This occurs if E_1, E_2, E_3 , and E_4 are satisfied. Because E_2 and E_3 are equally likely to be fulfilled by $i = 1$ or $i = 2$ and each is equally likely to occur, we can pick $i = 1$ and multiply by two. While it is easy to evaluate if any given game results in the social optimum being played, it is difficult to organize the integrals necessary to construct the probability of it occurring. To help visualize the integrals, we consider the graphical representation of a 2×2 game given in Fig. B.12. For a fixed a_{11}, b_{11} , we have

$$P(a_{11} + b_{11} \geq a_{21} + b_{21}, a_{22} + b_{22}, E_2, E_3 | a_{11}, b_{11}) = Q(a_{11}, b_{11}) \\ = \int_{-\infty}^{a_{11}} \left(\int_{-\infty}^{a_{11}+b_{11}-a_{21}} G(a_{11}, b_{11}, b_{21}) f_b(b_{21}) db_{21} \right) f_a(a_{21}) da_{21}$$

where

$$G(a_{11}, b_{11}, b_{21}) = \int_{-\infty}^{b_{21}} \left(\int_{-\infty}^{a_{11}+b_{11}-b_{22}} f_a(a_{22}) da_{22} \right) f_b(b_{22}) db_{22}$$

For 1,1 to be the social optimum and be played by player 2 if player 1 plays 1, it is necessary and sufficient that $b_{12} \leq b_{11}$ and $a_{11} + b_{11} \geq a_{12} + b_{12}$. Since

$$P(a_{11} + b_{11} \geq a_{12} + b_{12}, b_{11} \geq b_{12} | a_{11}, b_{11}) = H(a_{11}, b_{11})$$

$$H(a_{11}, b_{11}) = \int_{-\infty}^{b_{11}} \left(\int_{-\infty}^{a_{11}+b_{11}-b_{12}} f_a(a_{12}) da_{12} \right) f_b(b_{12}) db_{12}$$

is independent of $P(a_{11} + b_{11} \geq a_{21} + b_{21}, a_{22} + b_{22}, E_2 | a_{11}, b_{11})$, we then have

$$P(E_1, E_2, E_4 | a_{11}, b_{11}) \\ = P(a_{11} + b_{11} \geq a_{12} + b_{12}, b_{11} \geq b_{12} | a_{11}, b_{11}) \\ \cdot P(a_{11} + b_{11} \geq a_{21} + b_{21}, a_{22} + b_{22}, E_2 | a_{11}, b_{11}) \\ = Q(a_{11}, b_{11}) \cdot H(a_{11}, b_{11}) \\ = M(a_{11}, b_{11})$$

Since $P(E_1, E_2, E_3, E_4 | a_{11}, b_{11}) = P(E_1, E_2, E_3, \neg E_4 | a_{11}, b_{11})$, this tells us the probability that 1,1 is a Stackelberg equilibrium and the social optimum is

$$P = 2 \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} M(a_{11}, b_{11}) f_a(a_{11}) da_{11} f_b(b_{11}) db_{11}$$

Since all 4 entries in a 2×2 game are equally likely to be the social optimum and be played, this gives us

$$P(\text{S.O.}) = 8 \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} M(a_{11}, b_{11}) f_a(a_{11}) da_{11} f_b(b_{11}) db_{11}$$

□

We see from Theorem 9 that although we have a method to exactly compute the probabilities of the social optimum being played in a zero-trust game, it is already quite complex for even a 2×2 game. Further, introducing the δ_i values for either the incomplete or complete knowledge policies vastly complicates the geometry of the space over which our probability distributions are computed. Therefore, assuming that f_a and f_b have finite expectation and variance, it will be generally more cost effective to use sampling methods to estimate the probability of playing the social optimum as well as related properties due to the relative ease of solving leader-follower games. Before we move on to Section 5 to do exactly that, we will first briefly discuss the expectation of zero trust leader-follower games, as a metric against which to measure the effectiveness of our policies.

Theorem 10. The expected payoff of each player i in a $m_1 \times m_2 \times \dots \times m_n$ zero-trust n -player Stackelberg game where

each player i 's payoffs for each entry in the probability tensor are generated iid from a distribution A_i is

$$E[u_i | A_1, \dots, A_n] = m_i \int_{-\infty}^{\infty} x f_{A_i}(x) (F_{A_i}(x))^{m_i-1} dx$$

where f_{A_i} is the probability density function for A_i and F_{A_i} is the cumulative distribution function.

Proof. Consider player i 's choice: player 1 makes its choice based on all other players' responses, however player 1's payoff is independent of these players. Players 2 through $i-1$ have also made choices based on all later players responses, but again each of their payoffs is independent of these later players. Thus, after players 1 through $i-1$ have made their choices, player i chooses between m_i possible responses, each with an iid payoff and each independent of the other $n-1$ players' payoffs despite the fact that previous players factored i 's action into their choices, and i will factor later players into its choices. Player i will therefore be choosing the maximum of m_i iid variables and its payoff is distributed as the maximum of m_i samples from its payoff distribution A_i .

Let X be a random variable and let $X^{(k)}$ be the maximum of k iid samples of X . For its CDF we have $F_X^{(k)}(x) = (F_X(x))^k$ which gives us a PDF of $f_X^{(k)}(x) = k f_X(x) (F_X(x))^{k-1}$. Therefore, the expected value of $X^{(k)}$ is

$$E(X^{(k)}) = k \int_{-\infty}^{\infty} x f_X(x) (F_X(x))^{k-1} dx$$

which means the expected value of the game for player i is

$$E(u_i | A_1, \dots, A_n) = E(A_i^{(m_i)}) = m_i \int_{-\infty}^{\infty} x f_{A_i}(x) (F_{A_i}(x))^{m_i-1} dx.$$

□

Theorem 11. The social optimum of a $m_1 \times m_2 \times \dots \times m_n$ zero-trust n -player game with each player's payoffs generated iid from A_i is

$$\mathcal{M} \int_{-\infty}^{\infty} x f_X(x) (F_X(x))^{\mathcal{M}-1} dx$$

where X is distributed as $\sum_{i=1}^n A_i$ for independent A_i and $\mathcal{M} = \prod_{i=1}^n m_i$.

Proof. The social optimum is defined as the maximum value of $\sum_{i=1}^n u_i(\sigma)$ across all strategies $\sigma \in \Sigma$, and is always a pure strategy profile. There are $\mathcal{M}$ such pure-strategy σ , all iid. By the same reasoning as in the proof of Theorem 10, this means the expected value of the social optimum is

$$\mathcal{M} \int_{-\infty}^{\infty} x f_X(x) (F_X(x))^{\mathcal{M}-1} dx$$

for $X \sim \sum_{i=1}^n A_i$. □

Given the results of Theorems 10 and 11, we can conclude that the expected PoA of a zero-trust 2-player $m \times n$ leader-follower maximization game is

$$\frac{E[(A+B)^{(mn)}]}{E[A^{(m)}] + E[B^{(n)}]}$$

where $X^{(n)}$ is the maximum of n iid samples of a distribution X .

References

- Anderhub, V., Engelmann, D., & Gth, W. (2002). An experimental study of the repeated trust game with incomplete information. *Journal of Economic Behavior & Organization*, 48(2), 197–216.
- Aumann, R. J., & Shapley, L. S. (1994). *Long-term competition—a game-theoretic analysis* (pp. 1–15). New York, NY: Springer New York.
- Bapna, R., Qui, L., & Rice, S. C. (2017). Repeated interactions vs. social ties: Quantifying the economic value of trust, forgiveness, and reputation using a field experiment. *MIS Quarterly*, 41(3), 841–866.
- Benoit, J.-P., & Krishna, V. (1985). Finitely repeated games. *Econometrica*, 53(4), 905–922.

- Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1), 122–142.
- Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M., & Papaioannou, E. (2010). The impact of altruism on the efficiency of atomic congestion games. In M. Wirsing, M. Hofmann, & A. Rauschmayer (Eds.), *Trustworthy global computing* (pp. 172–188). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Chen, P.-A., Keijzer, B. D., Kempe, D., & Schäfer, G. (2014). Altruism and its impact on the price of anarchy. *ACM Trans. Econ. Comput.*, 2(4), 17:1–17:45.
- Chen, S.-P., & Wu, W.-Y. (2010). A systematic procedure to evaluate an automobile manufacturer distributor partnership. *European Journal of Operational Research*, 205(3), 687–698.
- Chen, X., & Deng, X. (2006). Settling the complexity of two-player nash equilibrium. In *Proceedings of the 47th annual IEEE symposium on foundations of computer science*. In FOCS 06 (pp. 261–272). USA: IEEE Computer Society.
- Choi, S., & Messinger, P. R. (2016). The role of fairness in competitive supply chain relationships: An experimental study. *European Journal of Operational Research*, 251(3), 798–813.
- Dal Bó, P., & Frchette, G. R. (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, 56(1), 60–114.
- Dasgupta, P. (2000). Trust as a commodity. *Trust: Making and breaking cooperative relations*, 4, 49–72.
- Du, S., Nie, T., Chu, C., & Yu, Y. (2014). Reciprocal supply chain with intention. *European Journal of Operational Research*, 239(2), 389–402.
- Engle-Warnick, J., & Slonim, R. L. (2004). The evolution of strategies in a repeated trust game. *Journal of Economic Behavior & Organization*, 55(4), 553–573. Trust and Trustworthiness
- Engle-Warnick, J., & Slonim, R. L. (2006a). Inferring repeated-game strategies from actions: evidence from trust game experiments. *Economic Theory*, 28(3), 603–632.
- Engle-Warnick, J., & Slonim, R. L. (2006b). Learning to trust in indefinitely repeated games. *Games and Economic Behavior*, 54(1), 95–114.
- Fudenberg, D., & Maskin, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3), 533–554.
- Gibbons, R. (2001). Trust in social structures: Hobbes and coase meet repeated games, *The Russell Sage Foundation Series on Trust*, volume 2 (pp. 332–353). New York: Russell Sage Foundation.
- Govindan, S., & Wilson, R. (2003). A global newton method to compute nash equilibria. *Journal of Economic Theory*, 110(1), 65–86. [https://doi.org/10.1016/S0022-0531\(03\)00005-X](https://doi.org/10.1016/S0022-0531(03)00005-X).
- Hamilton, W. D. (1963). The evolution of altruistic behavior. *The American Naturalist*, 97(896), 354–356.
- Hamilton, W. D. (1964a). The genetical evolution of social behaviour. i. *Journal of Theoretical Biology*, 7(1), 1–16.
- Hamilton, W. D. (1964b). The genetical evolution of social behaviour. ii. *Journal of Theoretical Biology*, 7, 17–52.
- Kakutani, S. (1941). A generalization of brouwers fixed point theorem. *Duke Math. J.*, 8(3), 457–459. <https://doi.org/10.1215/S0012-7094-41-00838-4>.
- Ledyard, J. O. (1994). Public Goods: A Survey of Experimental Research. *Public Economics*. University Library of Munich, Germany.
- Lemke, C. E. (1965). Bimatrix equilibrium points and mathematical programming. *Management Science*, 11(7), 681–689.
- Lemke, C. E., & Howson, J. T. (1964). Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2), 413–423.
- Leopold-Wildburger, U., Schuetze, J. H., & Lafer, A. (2002). Mccockereel measuring individual punishment and reciprocity in a simple value-laden dilemma game. *European Journal of Operational Research*, 140(2), 241–248.
- Morreale, A., Mittone, L., & Nigro, G. L. (2019). Risky choices in strategic environments: An experimental investigation of a real options game. *European Journal of Operational Research*, 279(1), 143–158. <https://doi.org/10.1016/j.ejor.2019.05.013>.
- Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1), 48–49.
- Niederhoff, J. A., & Kouvelis, P. (2016). Generous, spiteful, or profit maximizing suppliers in the wholesale price contract: A behavioral study. *European Journal of Operational Research*, 253(2), 372–382.
- Io Nigro, G., Morreale, A., Robba, S., & Roma, P. (2013). Biopharmaceutical alliance and competition: A real options games approach. *International Journal of Innovation Management*, 17(06), 1340023.
- Roughgarden, T. (2015). Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5), 32:1–32:42.
- Rubinstein, A. (1979). Equilibrium in supergames with the overtaking criterion. *Journal of Economic Theory*, 21(1), 1–9.
- Rubinstein, A. (1980). Strong perfect equilibrium in supergames. *International Journal of Game Theory*, 9(1), 1–12.
- Wilson, J. Q., & Kelling, G. L. (1982). Broken windows. *Atlantic monthly*, 249(3), 29–38.