A characterization of cooperation in dynamic network games

Caleb M. Koch*
Computational Social Science
Zurich, Switzerland
calebkoch@ethz.ch

Alain Rossier*
Department of Mathematics
Zurich, Switzerland
rossiera@student.ethz.ch

ABSTRACT

This paper explores conditions under which players cooperate in a dynamic network game. Historically, folk theorems have provided a speckled perspective by showing that there exists equilibria where players cooperate, do not cooperate, as well as a myriad of equilibria between these extremes. Our main contribution is identifying a necessary and sufficient equilibrium refinement such that, for all equilibria, all players cooperate in order to reach a strictly Pareto dominant graph. We base our results on a class of games that subsume forward-looking extensions of exchange economies with indivisible goods.

KEYWORDS

Networks, dynamic games, exchange economies, equilibrium refinement, stability, collusiveness, Pareto dominance

ACM Reference Format:
Caleb M. Koch and Alain Rossier. 2019. A characterization of cooperation in dynamic network games. In *The 14th Workshop on the Economics of Networks, Systems, and Computation, June 28, 2019, Phoenix, Arizona.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Can forward-looking, strategic, non-cooperative agents cooperate in order to achieve a mutually preferred outcome? Folk theorems have historically provided a mixed answer to this question. Broadly speaking, such theorems establish that *any* individually rational and feasible outcome is an equilibrium when players are *very* forward-looking [1, 5, 8]. As such, dynamic games with *very* forward-looking players have equilibria where players cooperate, never cooperate, as well as a myriad of equilibria between these extremes. Folk theorems are often viewed as a 'negative result' because it generally precludes the analysis of strategic behavior in such settings.

This paper presents an equilibrium refinement approach to shed new light on the posed question. Our main contribution is identifying a necessary and sufficient refinement such that, for all such equilibria, all players cooperate in order to reach a strictly Pareto dominant outcome. Our theorem is a 'positive result'. Specifically, this characterization reveals that a form of 'collusion' is necessary in order for non-cooperative players to cooperate notwithstanding a mutual incentive to do so. Perhaps somewhat surprisingly, the necessary and sufficient refinements that characterize these 'good' equilibria are considerably strong, highlighting that cooperation does not emerge naturally even from *very* forward-looking, strategic behavior.

To put our analysis into context, we study a class of dynamic network games that subsumes many applied economic models. In a companion paper [16], we show that the class subsumes dynamic extensions of many matching models [15] and exchange economies with indivisible goods [11]. For exchange economy models, our framework allows for general exchange processes (e.g., package models, [2]) and general preference structures (e.g., couple preferences, [17]). The class also subsumes dynamic extensions of commonly used myopic network models [13, 14] as well as the dynamic network game analyzed by [4].

Our equilibrium refinement is summarized in two conditions.

Refinement # 1: The first refinement, which we call *stability*, is closely related to trembling perfection [25]. We say that an action profile is unstable if a player does not form a link only because he/she anticipates the (respective) other to not form a link, notwithstanding a mutually strict benefit in doing so. Put another way, stability avoids the 'zero corner' of agents' action space.

Refinement # 2: The second refinement, which we call *collusiveness*, is related to the coordination issue described above.[1] However, whereas stability concerns players i and j who must *coordinate*, the second refinement requires players $k \neq i, j$ to *abstain* from accepting a second-best option in order to allow i and j to coordinate.

We prove that there exists an equilibrium that is stable and collusive in every dynamic network game (whether or not there exists a strict Pareto dominant outcome). Consequently, there always exists the potential that non-cooperative, strategic agents are able to cooperate, and our analysis highlights the necessary and sufficient barriers that must be overcome for this to occur.

Related literature. In the network literature, it is well-known that non-cooperative players do not necessarily cooperate in order to achieve collectively preferred outcomes. In myopic settings, it is often the case that agents lack the foresight to overcome locally inefficient outcomes in order to reach strictly preferred networks (see, e.g., [13] and [26]). While intuition suggests that introducing forward-looking strategies would attenuate this issue, Dutta et al. [4] present a negative example. Specifically, Dutta et al. (Theorem 2) demonstrate that forward-looking strategies, as predicted by Markov perfect equilibrium, exhibit additional complications that must be overcome for forward-looking players to 'agree' on

*Both authors contributed equally to this research.

[1] The collusion refinement is closely related to Herings et al. (2009) definition of a *"pairwise farsightedly stable set"*; see Section 4 for discussion.

NetEcon 2019, June 28, 2019, Phoenix, Arizona

ACM ISBN 978-1-4503-9999-9/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

outcomes that are strictly preferred by all players.[2] In this paper, we highlight a new problem that is not in their example (due to their link monotonicity assumption): players can prefer severing links to move away from strictly Pareto dominant outcomes in fear of others severing links and, thereby, creating unprofitable outcomes. In sum, forward-looking players generally will not reach strictly dominant outcomes, in the short- or long-run, notwithstanding a mutual strict preference in doing so and a discount factor tending to unity.

A different approach that has received attention is considering *perfectly farsighted* agents, such that agents make decisions only taking into account the long-run outcome [3, 20, 27]. In this extreme case, perfect farsightedness is a sufficient condition to ensure that non-cooperative players reach strictly Pareto dominant outcomes (see, e.g., Herings et al., 2009, Theorem 7). This paper addresses the less-extreme case of forward-looking considered by [4] and [9] in an attempt to identify conditions under which the positive results with perfectly farsighted players also hold with forward-looking players.

Structure of the paper. The paper proceeds as follows. In Section 2, we present the model. In Section 3, we present two equilibrium refinements, which form the basis of our main results in Section 4. We conclude in Section 5. All proofs can be found in the full companion paper [16].

2 DYNAMIC NETWORK GAMES

Players and networks. Let $\mathcal{N} = \{1, 2, \dots, n\}$ denote a set of ***players***. Players can form ***links*** between each other; we let ij represent a link between players i and j. No link between i and j is represented by $\emptyset$. It is useful to let $\mathcal{L}_{ij} = \{\emptyset, ij\}$ and $\mathcal{L}_i = \times_{j \neq i} \mathcal{L}_{ij}$.[3]

A ***graph*** is a collection of links between players. The set of possible graphs equals $\mathcal{G} \subseteq \{ij \mid i, j \in \mathcal{N} : i \neq j\}$. We denote a typical graph as $g \in \mathcal{G}$. For notation we let $g - ij = g \backslash \{ij\}$ represent the removal of link ij and $g + ij = g \cup \{ij\}$ represent the addition of link ij.

Utility. Player i's preference over graphs is represented by a ***utility*** $u_i : \mathcal{G} \to \mathbb{R}$. In general, i's utility depends on both personal links (e.g., ij) and links between other players (e.g., jk with $j, k \neq i$). The collection of all utilities is $u = (u_1, \dots, u_n) : \mathcal{G} \to \mathbb{R}^n$.

Actions. Players can change the graph topology in order to improve payoffs. The set of pure-strategies for player i can be represented as $\mathcal{A}_i^g = \mathcal{N}$ with the following interpretation: (i) if $ij \notin g$ then $j \in \mathcal{A}_i^g$ represents i trying to form a link with j; (ii) if $ij \in g$ then $j \in \mathcal{A}_i^g$ represents i severing the link with j; and (iii) $i \in \mathcal{A}_i^g$ represents i doing nothing.

A mixed-strategy at graph g is denoted as

$$\alpha_i(g) = (\alpha_{i1}(g), \dots, \alpha_{in}(g)) \in \Delta(\mathcal{A}_i^g).$$

Here, $\alpha_{ij}(g)$ is interpreted as the probability that i's realized action is $j \in \mathcal{A}_i^g$. An ***action***

$$\alpha_i = (\alpha_i(g))_{g \in \mathcal{G}} \in \times_{g \in \mathcal{G}} \Delta(\mathcal{A}_i^g) =: \mathscr{A}_i$$

is a mixed-strategy taken by player i at each graph $g \in \mathcal{G}$. We denote a collection of actions as $\alpha = (\alpha_1, \dots, \alpha_n) \in \times_{i \in \mathcal{N}} \mathscr{A}_i =: \mathscr{A}$.

Timing and how graphs change. Following [13] and [4], we suppose *one* player changes the graph topology at each time step $t = 0, 1, 2, \dots$ in the following way:

(1) At time step t, a player is selected according to a probability distribution $f(\cdot \mid g^t, \alpha)$, which may depend on the graph g^t and/or mixed-actions α.[4]
(2) If (e.g.) player i is selected, then he/she employs the mixed-strategy $\alpha_i(g^t) \in \Delta(\mathcal{A}_i^{g^t})$. The probability of the graph changing then depends on i's realized pure-strategy:
(2.a) Player i can sever any link unilaterally. For example, if $ij \in g$, then the probability that i chooses to sever its link with j is $\alpha_{ij}(g^t)$, in which case $g^{t+1} = g^t - ij$.
(2.b) Player i can form a link only if another player agrees to do so. For example, if $ij \notin g$, then the probability that i and j form a link is $\alpha_{ij}(g)\alpha_{ji}(g)$, in which case $g^{t+1} = g^t + ij$. $\alpha_{ij}(g)$ represents the probability that i proposes a link to j; $\alpha_{ji}(g)$ represents the probability that j accepts.[5]

Net-present valuations. We assume that players are forward-looking when strategically deciding whether to form or sever links. A player's ***net-present value*** of $g^t \in \mathcal{G}$, given a profile of strategies $\alpha \in \mathscr{A}$, equals

$$U_i(g^t, \alpha) := u_i(g^t) + \mathbb{E}\left[\sum_{s=1}^{\infty} \delta_i^s u_i(g^{t+s}) \,\middle|\, g^t, \alpha\right],$$

where $\delta_i \in (0, 1)$ is i's discount of future payoffs, and the expectation is formed by $P(\cdot \mid g, \alpha) \in \Delta(\mathcal{G})$, which is the conditional probability of moving from g^t to another graph induced by players' strategies α (the dependence on f is implied).

A ***dynamic network game*** is defined by the tuple

$$\Gamma := \{\mathcal{N}, \mathcal{G}, \mathscr{A}, u, f\}.$$

Example 2.1 (Bilateral free-trade agreement network [10]). An example of our model is a bilateral free-trade network game [10]. $\mathcal{N}$ represents countries. A link ij represents whether countries i and j have formed a bilateral free-trade agreement. $u_i : \mathcal{G} \to \mathbb{R}$ represents country i's payoff as a function of all countries' free-trade agreements. Depending on the free-trade agreement network $g \in \mathcal{G}$, $\alpha_{ij}(g)$ represents the probability that i ends the agreement with j / proposes a free-trade agreement with j. In the latter case, $\alpha_{ji}(g)$ represents the probability that j accepts i's free-trade agreement. Finally, $U_i(g^t, \alpha)$ represents country i's discounted flow of profits.

2.1 Equilibrium concept

We next introduce the equilibrium concept most commonly used to analyze dynamic games (see, e.g., [19], and [4]). Broadly speaking, an equilibrium action profile is one from which no player has a payoff incentive to deviate because each player plays a mutual best-reply action. To precisely define an equilibrium of a dynamic

[2] In contast to [4], we do not assume link monotonicity, anonymity, 'limited transfer', and 'increasing returns to link creation'.
[3] All results in the paper generalize to the case where $\mathcal{L}_{ij}$ is an arbitrary finite set; for details, see an earlier version of the paper.

[4] [13] and [4] assume that $f(i \mid g^t, \alpha) = 1/n$.
[5] The restriction to letting α_{ij} describe the probability of forming *and* accepting a link, rather than separating these two strategies, is for the sake of notation and not necessarily restrictive. All results would hold substantially but with more notation.

network game, we first must introduce what it means for an action to be *optimal*.

Definition 2.2. For player i, $\alpha_i^* \in \mathscr{A}_i$ is an ***optimal action*** with respect to opponents' action profile $\alpha_{-i} \in \mathscr{A}_{-i}$ if i's *ex ante* net-present value is maximized at every graph, i.e.,

$$U_i(g, \alpha_i^*, \alpha_{-i}) \geq U_i(g, \widehat{\alpha}_i, \alpha_{-i}) \quad \text{for all } g \in \mathcal{G} \text{ and } \widehat{\alpha}_i \in \mathscr{A}_i.$$

By definition, an optimal action takes into account all (probable) sequences of graphs and others' actions thereafter. An equilibrium action profile extends Definition 2.2 to all players.

Definition 2.3. A ***Markov perfect equilibrium (MPE)*** is an action profile $\alpha^* = (\alpha_1^*, \ldots, \alpha_n^*) \in \mathscr{A}$ such that all players play optimal actions, i.e., α_i^* is an optimal action with respect to α_{-i}^* for all $i \in \mathcal{N}$.

2.2 Assumption

To prove our main results, we require an assumption on the player selection process, f. Recall that $f(i \mid g, \alpha)$ is the probability that i is selected at graph g when players employ α. A common assumption is to let $f(i \mid g, \alpha) = 1/n$ [4, 13]. However, it is clear from our analysis that f can crucially determine whether or not players can cooperate to reach the strictly Pareto dominant outcome (see [16]).

In order to proceed, we impose the following assumption on the player selection process. The intuition is that players who are doing nothing—i.e. choose to stay at a current graph with probability one—are not selected by f. As such, only players actively trying to sever or propose a link are selected by f.

Assumption 1 (A-1). *A dynamic network game satisfies **A-1** if, for all $g \in \mathcal{G}$ and $\alpha \in \mathscr{A}$, $f(i \mid g, \alpha) > 0$ only if $\alpha_{ii}(g) < 1$.*

3 EQUILIBRIUM REFINEMENTS

In this section, we present two equilibrium refinements of MPE, which form the basis of our main results below.

3.1 Preliminary definitions

In what follows, we focus on dynamic network games that have a strictly Pareto dominant graph, which is defined as follows.

Definition 3.1. A graph g^S is ***strictly Pareto dominant*** if $u_i(g^S) > u_i(g)$ for all $i \in \mathcal{N}$ and all $g \neq g^S$.

We focus on games with a strictly Pareto dominant graph because it provides a baseline. In addition, it highlights that, even in such an extreme situation, it is nevertheless difficult for very forward-looking players to cooperate.

Our main result concerns when players can cooperate to reach a strictly Pareto dominant graph. To make our research question more precise, we must clarify what it means for players to 'reach' g^S. Denote P as the $|\mathcal{G}| \times |\mathcal{G}|$ matrix of transition probabilities induced by an action profile $\alpha \in \mathscr{A}$. As such, $(P^t)_{gg'}$ equals the probability of going from g to g' in exactly t steps.[6]

Definition 3.2. For a given action profile $\alpha \in \mathscr{A}$, a network g^S is ***strongly absorbing*** if, for all $g \in \mathcal{G}$, the Markov chain induced by α converges to g^S a.s., that is, $\lim_{t\to\infty}(P^t)_{gg^S} = 1$ for all $g \in \mathcal{G}$.

[6]To clarify, if $t = 0$ then P^0 is the identity matrix, and if $t = 1$ then $P_{gg'}$ is a matrix version of $P(g' \mid g, \cdot)$.

In other words, our measure for whether players are successful at attaining g^S is relatively weak: we simply ask if players, behaving according to a MPE, attain g^S *eventually* in infinite time irrespective of the starting graph. Notwithstanding our weak criterion, we show that MPE itself is far from sufficient to yield a positive answer.

It is also useful to utilize a somewhat weaker stability notion, *pairwise stability*, as in [13].

Definition 3.3. A graph $g \in \mathcal{G}$ is ***pairwise stable*** for a given action profile $\alpha \in \mathscr{A}$ if $P(g' \mid g, \alpha) = 0$ for all $g' \in \mathcal{G}\backslash\{g\}$.

In other words, once players reach a pairwise stable graph, they stay there indefinitely. Note that if a graph other than g^S is pairwise stable, then g^S is not strongly absorbing. It is clear that, for any pairwise stable graph $g \in \mathcal{G}$, $U_i(g) = \frac{1}{1-\delta_i} u_i(g)$.

Finally, we introduce a definition on which both refinement conditions below build. Let $\Sigma := \mathcal{N} \times \mathcal{G}$ denote the set of player-graph pairs. For some subset $S \subseteq \Sigma$ denote the corresponding profile of actions as $\alpha_S = (\alpha_j(g))_{(j,g)\in S}$. With some abuse of notation, we let $B_i(g, \alpha_{-S}, \alpha_S)$ denote i's set of best-replies to (α_{-S}, α_S), where $\alpha_{-S} \in \Delta(\mathcal{A}_{-S}) := \times_{(k,g')\in\Sigma\backslash S} \Delta(\mathcal{A}_k^{g'})$.

Definition 3.4. For some $\hat{\alpha} \in \mathscr{A}$ and $S \subset \Sigma$, $\hat{\alpha}_S$ is a ***robust subset of actions*** if, for all $(i, g) \in S$, $\hat{\alpha}_{S,i}(g) \in B_i(g, \alpha_{-S}, \hat{\alpha}_S)$ for all $\alpha_{-S} \in \Delta(\mathcal{A}_{-S})$.

Put another way, a robust subset of actions is a collection of actions, identified by $S \subset \Sigma$, such that each element is a best-reply in response to this set of actions identified by S, irrespective of the actions taken by player-graph pairs outside of S. We illustrate the definition of a robust subset of actions via examples below.

3.2 First refinement: Stability

The first refinement builds on the observation that MPE are in general not trembling perfect in the sense of [25]. If i perceives j as strictly preferring a network g^S, then i should anticipate j as placing positive probability weight on realizing g^S; as a result, i should in anticipation rationalize to also place probability weight on realizing this network if he, too, strictly prefers g^S. Yet, if j places zero probability weight on reaching g^S, then it is considered optimal by i to also place zero weight on linking with j, and vice versa. This argument does not hold if j or i's action space is perturbed by ε and $\alpha_{ij}(g^S) = 0$ was no longer a viable action. As such, an ε-tremble in the action profile can cause players to quickly adjust actions that were previously seen as optimal. This intuition underlies our next definition.

Definition 3.5 (Unstable triplet). For $\hat{\alpha} \in \mathscr{A}$ and some $S \subset \Sigma$, suppose $\hat{\alpha}_S$ is a robust subset of actions. Then for $(i, g), (j, g) \notin S$, the tuple (i, j, g) is an ***unstable triplet*** if the following holds for all $\alpha_{-S} \in \Delta(\mathcal{A}_{-S})$:

(i) $\exists\, \tilde{\alpha}_i(g) \in B_i(g, \alpha_{-S}, \hat{\alpha}_S)$ such that $\tilde{\alpha}_{ij}(g) = 0$ only if $\alpha_{ji}(g) = 0$,
(ii) $\exists\, \tilde{\alpha}_j(g) \in B_j(g, \alpha_{-S}, \hat{\alpha}_S)$ such that $\tilde{\alpha}_{ji}(g) = 0$ only if $\alpha_{ij}(g) = 0$.

Definition 3.6 ((Un-)stable action profiles). An action profile $\hat{\alpha} \in \mathscr{A}$ is ***unstable*** if there exists an unstable triplet (i, j, g) such that $\hat{\alpha}_{ij}(g) = 0$ or $\hat{\alpha}_{ji}(g) = 0$. An action profile is ***stable*** if it is not unstable.

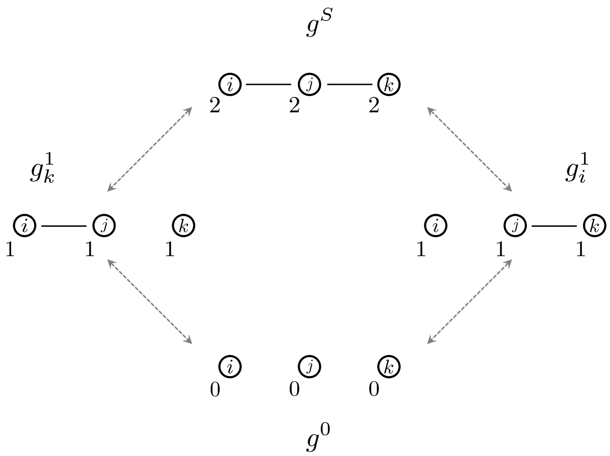

Figure 1: Example of stability.

The intuition for an unstable action profile is as follows. An unstable triplet requires that (e.g.) player i does not place probability weight to form a link with j only if i anticipates that j also does not try to form a link; the same must be true for j. An unstable action profile then requires that, indeed, neither player places probability weight on forming a link. Accordingly, we define stability as the negation of instability.

We illustrate stability with a simple example.

Example 3.7 (Stability — Figure 1). Consider players $\mathcal{N} = \{i, j, k\}$ who can form links between each other. We assume that j can form a link with i and/or k, while i and k cannot form a link. As such, possible network topologies consist of no links (which we denote as g^0), one link (g_k^1 or g_i^1), and two links (g^S). All players receive a utility equal to the number of links, as depicted in Figure 1. Therefore, the two-link graph is strictly Pareto dominant.

Intuitively, all players have the strategic incentive to place probability weight on reaching g^S from all graphs. We show that, in general, this is not the case.

Consider an action profile $\alpha^* \in \mathscr{A}$ such that $\alpha^*_{ii}(g^S) = \alpha^*_{jj}(g^S) = \alpha^*_{kk}(g^S) = 1$, i.e., the strictly Pareto dominant graph is pairwise stable. This implies that, letting $S := \{(l, g^S)\}_{l \in \mathcal{N}}$, α^*_S is a robust subset of actions. In addition, suppose that $\alpha^*_{ii}(g_k^1) = \alpha^*_{jj}(g_k^1) = \alpha^*_{kk}(g_k^1) = 1$, namely, players stay at g_k^1 with probability one and do not reach g^S despite being one link away. It is evident that (j, k, g_k^1) is an unstable triplet, and hence α^* is an unstable action profile.

We show α^* is indeed consistent with MPE behavior. Consider k's best-reply correspondence at g_k^1. If k places probability weight on linking with j, then k's net-present value would not change because the probability measure does not change, as j would not accept k's link proposal. Therefore, remaining at g_k^1 with the strategy $\alpha^*_{kk}(g_k^1) = 1$ and $\alpha^*_{kj}(g_k^1) = 0$ is in k's best-reply correspondence. Therefore, α^* is a MPE, unstable, and g^S is not strongly absorbing.

Note that α^* is not necessarily stable to ε-perturbations of the dynamic network game. Suppose that j places ε weight on linking with k (either because of an imperfect action or an ε-change in j's action space) and $(1 - \varepsilon)$ on remaining at g_k^1—note that this ε-deviation is in fact an optimal action. Given this new action profile, it is no longer optimal for k to place zero weight on reaching g^S. In contrast, it is uniquely optimal to set $\alpha_{kj}(g_k^1) = 1$ and $\alpha_{kk}(g_k^1) = 0$ irrespective of actions at all other graphs.

3.3 Second refinement: Collusiveness

The second refinement is related to stability. However, whereas stability addresses players i and j who must coordinate to *act*, the condition here is related to players $k \neq i, j$ who must *abstain* from acting in order to allow i and j to coordinate. The problem is that, if i and j fail to coordinate, other players accept a 'second best option' instead of waiting for i and j. We show later that MPE is not sufficient to preclude such behavior even for very forward-looking players.

To state the next definition, for any $\mathcal{K} \subseteq \mathcal{N}$ we let $\alpha_{\mathcal{K} \to g} := (\alpha_k(g))_{k \in \mathcal{K}}$ such that $\alpha_{kk}(g) = 1 \ \forall k \in \mathcal{K}$—that is, picking any player in $\mathcal{K}$ results in staying at g with probability one.

Definition 3.8. For $S \subset \Sigma$ and $\hat{\alpha} \in \mathscr{A}$, let $\hat{\alpha}_S$ be a robust subset of actions. Let $g \in \mathcal{G}$ and $\mathcal{K} \subset \mathcal{N}$ such that, for $S_{\mathcal{K}} := \bigcup_{k \in \mathcal{K}} \{(k, g)\}$, $S_{\mathcal{K}} \cap S = \emptyset$. Then we say that $(\mathcal{K}, S, g)$ is a ***collusive tuple*** if, for all $k \in \mathcal{K}$, $\exists \widetilde{\alpha}_k(g) \in B_k(g, \beta, \alpha_{\mathcal{K} \to g}, \hat{\alpha}_S)$ such that $\widetilde{\alpha}_{kk}(g) = 1$ for all $\beta \in \Delta(\mathcal{A}_{-(S \cup S_{\mathcal{K}})})$.

In other words, two conditions must be met in order for $(\mathcal{K}, S, g)$ to form a collusive tuple. First, S must form a robust subset of actions, which serves as the incentive for collusion. Second, at some $g \in \mathcal{G}$ there must exist a set of player $\mathcal{K}$ that would all strictly benefit to collude by setting $\alpha_{kk}(g) = 1$. This definition allows us to present our second refinement property.

Definition 3.9 ((Un-)collusive action profiles). An action profile $\hat{\alpha} \in \mathscr{A}$ is ***uncollusive*** if, for some collusive tuple $(\mathcal{K}, S, g)$, $\hat{\alpha}_{kk}(g) < 1$ for some $k \in \mathcal{K}$. An action profile is ***collusive*** if it is not uncollusive.[7]

Uncollusiveness extends Definition 3.8 by requiring that one player in $\mathcal{K}$ plays $\alpha_{kk}(g) < 1$, in spite of the incentive to collude.

We illustrate collusiveness with an example.

Example 3.10 (Collusiveness — Figure 2). Consider a dynamic network game consisting of players $\mathcal{N} = \{i, j, k.l\}$, as depicted in Figure 2. There exists four possible graphs, $\mathcal{G} = \{g^k, g^l, g^1, g^S\}$, where we assume g^S is strictly Pareto dominant. Possible actions at each graph is depicted in Figure 2. Reaching g^S from g^1 only requires a bilateral link between $\{i, j\}$, and reaching g^1 from g^k and g^l requires a bilateral link between $\{i, k\}$ and $\{j, k\}$, respectively.

We claim that, without the collusive property, stablility is not strong enough to ensure that g^S is strongly absorbing any MPE.

[7] The collusive property is closely related to the *"pairwise farsightedly stable set"* (PFSS) introduced by Herings et al. (2009, Definition 4). By condition (ia) of a PFSS, no player has an incentive to leave a PFSS for fear of arriving at another graph in the PFSS that is strictly worse off than the original graph. This relates to condition (ii) of uncollusiveness: trying to collude is deterred by the threat of $j \in \mathcal{I}$ acting ($\hat{\alpha}_{jj}(g) < 1$) and procuring unfavorable networks. Condition (ii) of a PFSS suggests that, broadly speaking, the PFSS is a (possible weak) subset of the recurrent class (see Koch and Rossier, 2019, for the definition and use of recurrent classes in our proof). Though these definitions are similar, the analysis in [12] takes a perspective closer to cooperative game theory while the analysis in this paper takes a non-cooperative game theory approach.

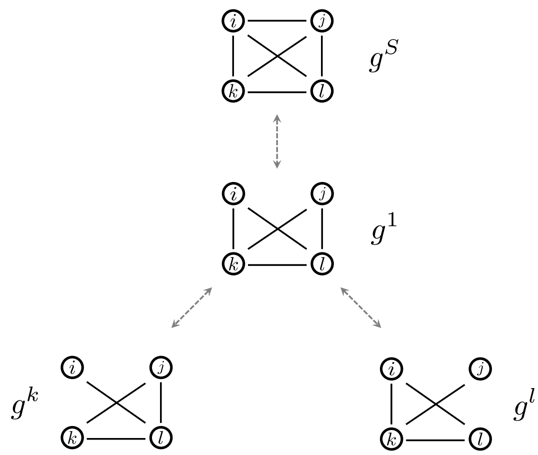

Figure 2: Example of collusiveness.

In fact, we claim that an uncollusive and stable MPE exists that renders g^k and g^l as pairwise stable outcomes even as $\delta \longrightarrow 1$.

Denote $\alpha^* \in \mathscr{A}$ as a stable MPE. By stability, $\alpha^*_{ij}(g^1) = \alpha^*_{ji}(g^1) = 1$. Suppose that $\alpha^*_{kk}(g^k) = \alpha^*_{ll}(g^l) = 1$ (k and l stay at g^k and g^l with probability one) and $\alpha^*_{ki}(g^1) = \alpha^*_{lj}(g^1) = 1$ (k and l leave g^1 with probability one).

We must show that the action profile as described is uncollusive. Firstly, it can easily be shown that, for $S = \{(i, g^1), (j, g^1)\}$, α^*_S is a robust subset of actions. Secondly, if k and l set $\alpha'_{kk}(g^1) = \alpha'_{ll}(g^1) = 1$ then g^S is reached from g^1 with probability one, which would be strictly preferred by both k and l; therefore, with $\mathcal{K} = \{k, l\}$, $(\mathcal{K}, S, g^1)$ is a collusive tuple. Finally, both $\alpha^*_{kk}(g^1) = 0 < 1$ and $\alpha^*_{ll}(g^1) = 0 < 1$. Therefore, α^* is uncollusive.

Next, we must show that k's net-present value of g^k is strictly greater than g^1, which would render g^k as pairwise stable and g^S as not strongly absorbing. In other words (omitting α^* from notation momentarily),

$$\begin{aligned} U_k(g^k) &= \frac{1}{1-\delta_k} u_k(g^k) \\ &> U_k(g^1) \\ &= u_k(g^1) + \delta_k \cdot \frac{1}{4}\left(2U_k(g^S) + U_k(g^k) + U_k(g^l)\right). \end{aligned} \tag{1}$$

If (1) holds, then we show in the companion paper that setting $\alpha_{kk}(g^k) = 1$ for k is a unique best-reply action ([16], Corollary 2). By symmetry, we can show that the same holds for l, which establishes that the action profile is indeed a MPE.

Player l's action, rendering g^l as pairwise stable, implies that $U_k(g^l) = \frac{1}{1-\delta_k} u_k(g^l)$. Using Corollary 2 in the companion paper again, strict Pareto dominance implies that $U_k(g^S) = \frac{1}{1-\delta_k} u_k(g^S)$. These observations and re-arranging terms in (1) yield

$$(4 - \delta_k) u_k(g^k) > 4(1 - \delta_k) u_k(g^1) + \delta_k \cdot \left(2u_k(g^S) + u_k(g^l)\right)$$

and letting $\delta_k \longrightarrow 1$ simplifies the expression to

$$3u_k(g^k) > 2u_k(g^S) + u_k(g^l).$$

From this expression, it is clear that the inequality can be satisfied for suitably selected utility values. The intuition is that l's threat of moving from g^1 to g^l with $\alpha^*_{lj}(g^1) = 1$ incentivizes k to not wait for i or j to be selected and, instead, place probability weight on the second-best option of g^k. Therefore, without the collusive property, there exists a dynamic network game Γ with a stable MPE such that g^S is not realized in the long-run as $\delta \longrightarrow 1$.

The collusive property precludes l's threat-based action. Clearly, if $\alpha_{ll}(g^1) = 1$ then k's unique best-reply action at g^1 is to also place all probability weight on realizing g^S, i.e., $\alpha_{kk}(g^1) = 1$. The same holding true for l, implying that both players strictly prefer colluding by setting $\alpha_{kk}(g^1) = \alpha_{ll}(g^1) = 1$.

3.4 Summary

In sum, stability is a refinement that ensures that players are coordinating to form mutually beneficial links. Collusiveness, on the other hand, is a refinement that focuses on making sure that players *not* forming links help players that are formings links toward mutually preferred outcomes. As shown below, these refinements together are both necessary and sufficient for very forward-looking, non-cooperative players to cooperate and reach a strictly Pareto dominant graph.

4 MAIN RESULTS

In this section, we present our main results. These results take two forms. First, we prove that stable and collusive MPE exist in any dynamic network game. Second, we prove that stable and collusive as equilibrium refinements are necessary and sufficient for ensuring that all such MPE render g^S as strongly absorbing.

All proofs can be found in the companion paper [16].

4.1 Characterization of cooperation

We now present our main results.

4.1.1 Existence of a stable and collusive MPE. In the following theorem, we establish that there exists a MPE that satisfies stability and collusiveness in every dynamic network game.

Theorem 4.1. *In all dynamic network games, there exists a (stationary) MPE that satisfies stability and collusiveness.*

The proof builds on previous equilbrium existence results of finite dynamic games from [7] and [4], and we demonstrate how to construct a stable and collusive MPE. In particular, the existence of a MPE is ensured by Kakutani fixed-point theorem (see [16]) Suppose, contrary to our assertion, that there exists no equilibrium that is stable. Then by definition, there exists a robust subset of actions. We construct a reduced action space where (i) all robust subset of actions are fixed, (ii) the 'unstable' action space region is reduced by ε-amount, and (iii) the remaining action space is unchanged. We construct the region in such a way that the Kakutani fixed-point theorem again applies. If there exists no stable MPE, we apply the same argument. We can repeat only finitely many times until we eventually arrive at a stable MPE, which contradicts our initial hypothesis; therefore, there exists a stable MPE. The argument for the existence of a stable *and* collusive MPE follows similarly.

4.1.2 Characterization of cooperation in dynamic network games. Before we present our main result, some notation is required. Define $\mathcal{E}^{MPE}(\Gamma_\delta)$ as the set of MPE of a dynamic network game Γ_δ, where $\delta = (\delta_1, \delta_2, \ldots, \delta_n)$ represents players' discount factors. We compare players' discount factors with the standard product order such that $\overline{\delta} \geq \underline{\delta}$ if and only if $\overline{\delta}_i \geq \underline{\delta_i}$ for all $i = 1, 2, \ldots, n$.

Let ρ denote an equilibrium refinement property, e.g., stability and/or collusiveness. We then denote $\mathcal{E}^{MPE}(\cdot, \rho) \subseteq \mathcal{E}^{MPE}(\cdot)$ as the set of all MPE that satisfy property ρ.

Theorem 4.2. *There exists a $\overline{\delta} \in (0,1)^n$ such that, for all $\delta \geq \overline{\delta}$, the following statements are equivalent:*

(1) For all dynamic network games Γ_δ satisfying A-1 with a strictly Pareto dominant graph g^S, g^S is strongly absorbing for all $\alpha^* \in \mathcal{E}^{MPE}(\Gamma_\delta, \rho)$;
(2) Refinement ρ represents stability and collusiveness.

In other words, stability and collusiveness are necessary and sufficient conditions to ensure that, for the *global* class of dynamic network games, players attain collectively preferred outcomes. It is evident in the proof [16] that collusiveness is a necessary condition only for games in which players have strong incentives to miscoordinate. In a particular dynamic network game, the collusive refinement may be vacuous. With the following corollary we show that stability, on the other hand, is a necessary refinement that applies for all dynamic networks games.

Corollary 4.3. *Let Γ denote a dynamic network game with a strictly Pareto dominant graph, g^S. Then there exists a $\overline{\delta} \in (0,1)^n$ such that, for all $\delta \geq \overline{\delta}$, g^S is strongly absorbing for all $\alpha^* \in \mathcal{E}^{MPE}(\Gamma_\delta, \rho)$ only if ρ implies stability.*

5 DISCUSSION

In this paper, we explore conditions under which *very* forward-looking, strategic players cooperate in a dynamic network game. Historically, folk theorems have provided a speckled (and in many ways a negative) perspective on whether non-cooperative players can cooperate in dynamic settings. We present an equilibrium refinement approach to go beyond folk theorems. Our main contribution is identifying a necessary and sufficient equilibrium refinement that guarantees the selection of a strictly Pareto dominant graph (if one exists) in the long run as the discount factor tends to one. The refinements we identify are considerably strong, highlighting that cooperation does not emerge naturally even from *very* forward-looking, strategic behavior.

We have made several assumptions to benchmark our main results. In particular, we assumed common knowledge, focused on games with a strictly Pareto dominant graph, and on a particular link formation dynamic. These assumptions clarify how difficult it is for non-cooperative, very forward-looking agents to cooperate. One avenue of future work is revisiting these assumptions and, more generally, subjecting our main results to a robustness check.

There exists several avenues for future work. One possible avenue is investigating the role of information in our results. Throughout, we assumed complete information (e.g., players knew the graph structure at all states and there exists no private information), and it remains an open question whether environments with asymmetric information [6, 18] or 'no information' [21, 22] support or hinder forward-looking, strategic players from attaining efficient outcomes. Another possible avenue is considering this problem from a mechanism design perspective. To this end, our necessary and sufficient refinement provides a road map for developing mechanisms that facilitate forward-looking, strategic players to attain collectively preferred outcomes [23, 24].

REFERENCES

[1] Robert J Aumann and Lloyd S Shapley. 1994. Long-term competition—A game-theoretic analysis. In *Essays in Game Theory*, N Megiddo (Ed.). Springer, New York, NY, 1–15.
[2] Sushil Bikhchandani and Joseph M Ostroy. 2002. The package assignment model. *Journal of Economic Theory* 107, 2 (2002), 377–406.
[3] Michael Suk-Young Chwe. 1994. Farsighted coalitional stability. *Journal of Economic Theory* 63, 2 (1994), 299–325.
[4] Bhaskar Dutta, Sayantan Ghosal, and Debraj Ray. 2005. Farsighted network formation. *Journal of Economic Theory* 122, 2 (2005), 143–164.
[5] Prajit K Dutta. 1995. A folk theorem for stochastic games. *Journal of Economic Theory* 66, 1 (1995), 1–32.
[6] Chaim Fershtman and Ariel Pakes. 2012. Dynamic games with asymmetric information: A framework for empirical work. *Quarterly Journal of Economics* 127, 4 (2012), 1611–1661.
[7] Arlington M. Fink. 1964. Equilibrium in a stochastic n-person game. *Journal of Science of the Hiroshima University, Series A-I* 28, 1 (1964), 89–93.
[8] James W Friedman. 1971. A non-cooperative equilibrium for supergames. *Review of Economic Studies* 38, 1 (1971), 1–12.
[9] Douglas M Gale and Shachar Kariv. 2007. Financial networks. *American Economic Review* 97, 2 (2007), 99–103.
[10] Sanjeev Goyal and Sumit Joshi. 2006. Bilateralism and free trade. *International Economic Review* 47, 3 (2006), 749–778.
[11] Faruk Gul and Ennio Stacchetti. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87, 1 (1999), 95–124.
[12] P Jean-Jacques Herings, Ana Mauleon, and Vincent Vannetelbosch. 2009. Farsightedly stable networks. *Games and Economic Behavior* 67, 2 (2009), 526–541.
[13] Matthew O Jackson and Alison Watts. 2002. The evolution of social and economic networks. *Journal of Economic Theory* 106, 2 (2002), 265–295.
[14] Matthew O Jackson and Asher Wolinsky. 1996. A strategic model of social and economic networks. *Journal of Economic Theory* 71, 1 (1996), 44–74.
[15] Alexander S Kelso and Vincent P Crawford. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50, 6 (1982), 1483–1504.
[16] Caleb M. Koch and Alain Rossier. 2019. A characterization of cooperation in dynamic network games. *Mimeo: https://polybox.ethz.ch/index.php/s/7maRGODLSjPK2YO* (2019).
[17] Fuhito Kojima, Parag A Pathak, and Alvin E Roth. 2013. Matching with couples: Stability and incentives in large markets. *Quarterly Journal of Economics* 128, 4 (2013), 1585–1632.
[18] Qingmin Liu, George Mailath, Andrew Postelwaite, and Larry Samuelson. 2014. Stable matching with incomplete information. *Econometrica* 82, 2 (2014), 541–587.
[19] Eric Maskin and Jean Tirole. 2001. Markov perfect equilibrium: I. Observable actions. *Journal of Economic Theory* 100, 2 (2001), 191–219.
[20] Ana Mauleon and Vincent Vannetelbosch. 2004. Farsightedness and cautiousness in coalition formation games with positive spillovers. *Theory and Decision* 56, 3 (2004), 291–324.
[21] Heinrich H Nax and Bary SR Pradelski. 2015. Evolutionary dynamics and equitable core selection in assignment games. *International Journal of Game Theory* 44, 4 (2015), 1–30.
[22] Heinrich H Nax and Bary SR Pradelski. 2016. Core stability and core selection in a decentralized labor matching market. *Games* 7, 10 (2016), 1–16.
[23] Szilvia Pápai. 2007. Exchange in a general market with indivisible goods. *Journal of Economic Theory* 132, 1 (2007), 208–235.
[24] Alessandro Pavan, Ilya Segal, and Juuso Toikka. 2014. Dynamic mechanism design: A myersonian approach. *Econometrica* 82, 2 (2014), 601–653.
[25] Reinhard Selten. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4, 1 (1975), 25–55.
[26] Olivier Tercieux and Vincent Vannetelbosch. 2006. A characterization of stochastically stable networks. *International Journal of Game Theory* 34, 3 (2006), 351–369.
[27] Licun Xue. 1998. Coalitional stability under perfect foresight. *Economic Theory* 11, 3 (1998), 603–627.