

Temporal Network Creation Games

Davide Bilò¹, Sarel Cohen², Tobias Friedrich², Hans Gawendowicz²,
Nicolas Klodt², Pascal Lenzner² and George Skretas²

¹ University of L'Aquila, Italy

²Hasso Plattner Institute Potsdam, Germany

¹davide.bilo@univaq.it, ²{firstname.lastname}@hpi.de

Abstract

Most networks are not static objects, but instead they change over time. This observation has sparked rigorous research on temporal graphs within the last years. In temporal graphs, we have a fixed set of nodes and the connections between them are only available at certain time steps. This gives rise to a plethora of algorithmic problems on such graphs, most prominently the problem of finding temporal spanners, i.e., the computation of subgraphs that guarantee all pairs reachability via temporal paths. To the best of our knowledge, only centralized approaches for the solution of this problem are known. However, many real-world networks are not shaped by a central designer but instead they emerge and evolve by the interaction of many strategic agents. This observation is the driving force of the recent intensive research on game-theoretic network formation models.

In this work we bring together these two recent research directions: temporal graphs and game-theoretic network formation. As a first step into this new realm, we focus on a simplified setting where a complete temporal host graph is given and the agents, corresponding to its nodes, selfishly create incident edges to ensure that they can reach all other nodes via temporal paths in the created network. This yields temporal spanners as equilibria of our game. We prove results on the convergence to and the existence of equilibrium networks, on the complexity of finding best agent strategies, and on the quality of the equilibria. By taking these first important steps, we uncover challenging open problems that call for an in-depth exploration of the creation of temporal graphs by strategic agents.

1 Introduction

Networks are omnipresent in everyday life. They range from abstract constructs, such as (online) social networks, to essential infrastructure, such as transportation networks and power grids. Given their ubiquity and importance, rigorous research has been conducted to better understand real-world networks. Researchers strive to evaluate the behavior of networks, their

structural properties and the processes that drive their formation. Over the years, the research community has realized that more varied and complex models are required to capture the intricacies that govern a network's attributes. Two such intricacies are:

- (i) many networks are dynamic in nature, i.e. the nodes and/or the connections of the nodes change over time;
- (ii) the formation of many networks is driven by many individual and selfish agents without central coordination.

Due to the complexity that each of these settings introduces, researchers so far considered only one of the above assumptions that many real-world networks naturally exhibit.

However, in many real-world settings both (i) and (ii) apply. For example, consider the problem of scheduling meetings in a large institution where employees are interested in disseminating information to all their colleagues. For this, they can schedule meetings with others at different time slots depending on their availability. Meetings with multiple individuals at the same time slot enables the spread of information to all participants. Naturally, the goal is to minimize the number of meetings needed to inform everyone.

Our goal is for this paper to be the inaugural effort in combining dynamic networks with a game-theoretic analysis in order to better capture the formation of real-world networks.

1.1 Our Approach

We initiate the study of dynamic networks from a game-theoretic perspective by combining one of the earliest and very influential strategic network formation models for static networks, the non-cooperative network formation model by Bala and Goyal [2000], with the seminal temporal graph model of Kempe *et al.* [2002].

In the network formation model by Bala and Goyal [2000], the agents are nodes of a network and they strategically create costly incident links to maximize the number of nodes they can reach either directly or via a sequence of hops in the network. The temporal graph model of Kempe *et al.* [2002] assumes that an edge-labeled graph is given, where the labels indicate the time step where the respective edge is available. By combining the features of these two models, we assume that a temporal graph serves as the host graph for our network formation game. Agents correspond to its nodes and can create costly incident edges having the time labels specified by the host graph. Most importantly, instead of using standard

reachability defined as the existence of a path between two nodes, we employ the concept of temporal reachability, where some node u can reach a node v if a temporal path, i.e., a path with monotonically increasing edge labels, exists.

As a first step in this line of research, we consider a restricted version, where the underlying temporal host graph is a clique, all edges have unit cost, and the objective of each agent is to create as few edges as possible to ensure the existence of a temporal path from itself to every other node of the graph. Although being the simplest variant of our framework, this setting has the striking feature that equilibrium states of our game correspond to temporal spanners of the underlying temporal host graph. Thus, our model captures the decentralized creation of a temporal spanner by selfish agents. To the best of our knowledge, so far only centralized approaches exist for this prominent algorithmic problem.

We emphasize that our framework can be generalized to much more complex settings. In particular, and similarly to (recent variants of) the well-studied Network Creation Game by Fabrikant *et al.* [2003], more than temporal reachability could be studied in future work. For example, the existence of short temporal paths, additional robustness guarantees, and more complicated edge cost functions.

1.2 Our Contribution

We explore the formation of temporal spanners by strategic agents via studying the Temporal Reachability Network Creation Game, whose equilibrium networks must be temporal spanners. Besides this being the first decentralized approach for computing temporal spanners, the entailed equilibria must be stable with respect to local changes of the involved nodes.

Although we show that computing a best response strategy is NP-hard even if the host graph has a lifetime $t = 2$, we nonetheless show for this case that equilibria exist and that they can be computed efficiently. The existence of equilibria remains a challenging open problem for $t \geq 3$. However, we show that in this case, deciding if a given strategy profile is an equilibrium is NP-hard. This is remarkable, as similar questions are still open for most other game-theoretic network creation models. Also, in contrast to the classical Network Formation Game on static graphs [Bala and Goyal, 2000], this shows that incorporating temporal graphs yields a computationally much harder model.

As our main contribution, we provide non-trivial structural properties of equilibrium networks and we exploit them to prove bounds on the Price of Anarchy (PoA), i.e., on the quality of the obtained temporal spanners. Low bounds on the PoA imply that these equilibrium spanners are close to optimal. Regarding this, we give an upper bound of $\mathcal{O}(\sqrt{n})$ on the PoA and provide a lower bound of $\Omega(\log n)$.

Moreover, driven by the hardness of computing a best response strategy, we also investigate Greedy Equilibria (GE), that rely on very simple strategy changes. We connect them to Nash Equilibria by showing that the PoA with regard to Greedy Equilibria is at most a $\mathcal{O}(\log n)$ factor larger than the PoA with regard to Nash Equilibria. This shows that not much is lost by focusing on GEs.

All omitted details can be found in [Bilò *et al.*, 2023].

1.3 Related Work

The formation of networks by strategic agents has been studied intensively within the last decades. One of the earliest models is also closest to our work. In the Network Formation Game by Bala and Goyal [2000] selfish agents buy incident edges and their utility is a function that increases with the number of agents they can reach and it decreases with the number of edges bought. Most relevant for us is the version where undirected edges are formed. For this the authors prove that equilibria always exist and that they are either stars or empty graphs. Moreover, improving response dynamics quickly converge to such states. Also, for this model computing a best response strategy and deciding if a given state is in equilibrium can be done efficiently.

The network formation game was extended to a setting with attacks on the formed network [Goyal *et al.*, 2016]. There, the objective is post-attack reachability. This variant is more complex, but best response strategies can still be computed efficiently [Friedrich *et al.*, 2017]. Recently, a variant with probabilistic attack was studied [Chen *et al.*, 2019]. Also related are Topology Control Game [Eidenbenz *et al.*, 2006], where the agents are points in the plane and edge costs are proportional to the Euclidean distance among the endpoints. Similar in spirit is the model by Gulyás *et al.* [2015], but there the agents are points in hyperbolic space using greedy routing. Also models exist where the agents aim for creating a robust network, i.e., communication in the network should rely on more than a single path [Meirom *et al.*, 2015; Chauhan *et al.*, 2016; Echzell *et al.*, 2020].

Besides network formation games with reachability objective, even more model variants exist, where shortest path distances play a prominent role. Starting with the Network Creation Game [Fabrikant *et al.*, 2003], that is based on the even older Connections Game [Jackson and Wolinsky, 1996], researchers have focused on utility functions that depend on the distances of the respective agent to the other agents in the formed network. For this, variants exist that involve cooperation [Corbo and Parkes, 2005; Andelman *et al.*, 2009], locality [Bilò *et al.*, 2016; Cord-Landwehr and Lenzner, 2015], non-uniform edge prices [Chauhan *et al.*, 2017; Bilò *et al.*, 2019], and most recently, social networks [Bilò *et al.*, 2021; Friedrich *et al.*, 2022; Bullinger *et al.*, 2022].

For most of these variants, computing a best response strategy is NP-hard, improving response dynamics are not guaranteed to converge, and the hardness of deciding equilibria is open. Moreover, for the original Network Creation Game [Fabrikant *et al.*, 2003] equilibria always exist and the Price of Anarchy (PoA) is known to be constant for almost the full parameter range of the model. Similar results hold for the other models, with some notable exceptions, e.g., the geometric version has a high PoA [Bilò *et al.*, 2019].

To the best of our knowledge, no game-theoretic network formation model involving temporal graphs has been studied. However, starting from the work by Kempe *et al.* [2002], a lot of research has been devoted to algorithmic problems on temporal graphs, in particular to temporal spanners. Relevant for us, it has been shown that temporal cliques admit sparse temporal spanners [Casteigts *et al.*, 2021] and also sparse spanners with low stretch are possible [Bilò *et al.*, 2022]. In con-




Figure 1: The left shows a temporal graph with $n = 4$ nodes and lifetime $t = 3$. The sequence v_1, v_3, v_4, v_2 (red) is not a temporal path since its labels are not monotonically increasing. On the other hand, v_3, v_2, v_1, v_4 (blue) is a temporal path from v_3 to v_4 . The right shows the graph from the left as the host graph H and the graph $G(\mathbf{s})$ (not dashed) formed by the strategies of the agents. Here, v_1 plays greedy best response, since neither adding the edge (v_1, v_2) nor removing one of the edges (v_1, v_3) or (v_1, v_4) decreases its cost. However, v_1 does not play best response since removing (v_1, v_3) , (v_1, v_4) and buying (v_1, v_2) is an improving move.

trast, on non-complete temporal graphs spanners can be very dense [Axiotis and Fotakis, 2016]. The same holds true if strict temporal paths are considered, i.e., if the labels on the edges of a path must strictly increase [Kempe *et al.*, 2002]. Closely related to the reachability problem, [Klobas *et al.*, 2022] study the problem of finding the minimum number of labels required to achieve temporal connectivity in a graph.

Spanners on static graphs are a classical topic, see [Ahmed *et al.*, 2020] for a recent survey. For spanners in geometric settings, see [Narasimhan and Smid, 2007].

1.4 Model and Notation

Before stating our game-theoretic model, we will first introduce temporal graphs and temporal spanners.

Temporal Graphs and Spanners

A *temporal graph* $G = (V_G, E_G, \lambda_G)$ is an undirected labeled graph, where $\lambda_G: E_G \rightarrow \mathbb{N}$, assigns a label to each edge. The edge labels of G model at which point in time an edge is available. For simplicity, we assume that all labels are consecutive starting with 1. Formally, $\bigcup_{e \in E_G} \lambda_G(e) = \{1, 2, \dots, t\}$, for some $t \in \mathbb{N}$, where t is called the *lifetime*. Note that, for simplicity, we assume that every edge has a single time label, i.e., every edge is available only at a particular single time step. All results in this paper also hold if we extend the model to allow multiple labels per edge. As long as the graph G is clear from context, we might omit the subscripts and write V, E , and λ instead of V_G, E_G , and λ_G .

A (simple) *temporal path* in G is a (simple) path in G with monotonically increasing edge labels. Formally, it is a sequence of (distinct) nodes $v_1, \dots, v_i \in V_G$, that forms a (simple) path in G , i.e., for $1 \leq j \leq i - 1$ we have $e_j = \{v_j, v_{j+1}\} \in E_G$, where for all $1 \leq j \leq i - 2$ we have $\lambda_G(e_j) \leq \lambda_G(e_{j+1})$. Note, that the labels do not have to increase strictly since we assume zero edge traversal time.

For two nodes $u, v \in V_G$, we say that u can *reach* v in the temporal graph G if and only if there is a temporal path from u to v in G . We define $R_G(u)$ as the set of nodes that u can reach in G . Note that $u \in R_G(u)$, since every node can trivially reach itself via a temporal path of length 0. If every node can reach every other node, we say that G is *temporally*

connected. With this we define a *temporal spanner* of G as any temporally connected subgraph G' of G with $V_{G'} = V_G$ and $E_{G'} \subseteq E_G$. If no edge can be removed from G' while keeping the temporal spanner property, we call G' a *minimal temporal spanner* of G . If G' has at most as many edges as any other temporal spanner of G , we call G' a *minimum temporal spanner* of G .

The Temporal Reachability Network Creation Game

Now we define our game-theoretic network creation model, called the *Temporal Reachability Network Creation Game (TRNCG)*. Let $H = (V_H, E_H, \lambda_H)$ be a given complete temporal graph that serves as the *host graph* of our game. We assume that every node $v \in V_H$ corresponds to a strategic agent and let $|V_H| = n$ denote the number of agents.

We assume that agents play strategies, where a *strategy* S_v of some agent v is defined as $S_v \subseteq V_H \setminus \{v\}$, i.e., the strategy specifies to which other agents agent v wants to create an edge. The strategies of all agents together form the *strategy profile* $\mathbf{s} = \bigcup_{v \in V_H} \{(v, S_v)\}$.

A strategy profile \mathbf{s} defines the *created directed temporal graph* $G(\mathbf{s}) = (V_{G(\mathbf{s})}, E_{G(\mathbf{s})}, \lambda_{G(\mathbf{s})})$, with

$$\begin{aligned} V_{G(\mathbf{s})} &= V_H, \\ E_{G(\mathbf{s})} &= \{(u, v) \mid u, v \in V_H \wedge v \in S_u\}, \end{aligned}$$

where $\lambda_{G(\mathbf{s})}$ is the labeling λ_H restricted to the edge set $E_{G(\mathbf{s})}$, with $\lambda_{G(\mathbf{s})}((u, v)) = \lambda_{G(\mathbf{s})}((v, u)) = \lambda_H(\{u, v\})$. Thus, $G(\mathbf{s})$ is a directed subgraph of the complete host graph H that contains the union of all edges that are created by the agents, i.e., for every edge in $G(\mathbf{s})$ there is exactly one agent that wants to create it. We use directed edges to encode the owner of the edge, where edges are always directed away from their owner. For reachability, these edge directions will be ignored.

Agents choose their respective strategy to minimize their individual cost, where the cost of agent v in the created directed temporal graph $G(\mathbf{s})$ is defined as

$$c_H(v, \mathbf{s}) = |S_v| + K \cdot |V_H \setminus R_{G^{\leftrightarrow}(\mathbf{s})}(v)|,$$

where $K > 1$ is a large constant and $G^{\leftrightarrow}(\mathbf{s})$ is the undirected version of $G(\mathbf{s})$, i.e., $V_{G^{\leftrightarrow}(\mathbf{s})} = V_{G(\mathbf{s})}$ and $E_{G^{\leftrightarrow}(\mathbf{s})} = \{\{u, v\} \mid u \in S_v \vee v \in S_u\}$ with $\lambda_{G^{\leftrightarrow}(\mathbf{s})}(\{u, v\}) = \lambda_H(\{u, v\})$, for all edges $\{u, v\} \in E_{G^{\leftrightarrow}(\mathbf{s})}$. Thus, in the created temporal graph $G(\mathbf{s})$, agent v incurs a cost of one unit for each edge it creates and a penalty of K for each agent it cannot reach via a temporal path that ignores edge directions. Hence, agents aim to create as few edges as possible while still maintaining undirected temporal reachability.

Let $\mathbf{s}_{-v} = \mathbf{s} \setminus (v, S_v)$ denote the set of strategies of all agents other than agent v . Now consider that v changes its strategy S_v to S'_v . The resulting strategy profile is $\mathbf{s}_{-v} \cup \{(v, S'_v)\}$ which we will abbreviate as $\mathbf{s}_{-v} \cup S'_v$. We say that agent v 's strategy change from S_v to S'_v is an *improving move*, if it yields strictly less cost for agent v , i.e., if $c_H(v, \mathbf{s}_{-v} \cup S'_v) < c_H(v, \mathbf{s})$. If we additionally restrict

the strategy change to a single addition or deletion¹, we call it a *greedy improving move*. For an example, see Figure 1. If there is no improving move or greedy improving move of v for s , we call S_v a *best response* or *greedy best response*, respectively.

Given this definition, we can now define our solution concepts. For a given host graph H we say that the strategy profile s is in *Pure Nash Equilibrium (NE)* [Nash, 1950], if no agent has an improving move. We say s is in *Greedy Equilibrium (GE)* [Lenzner, 2012] if no agent has a greedy improving move. Note, that every greedy improving move is also an improving move and therefore every NE is also in GE.

Since we have a bijection between s and $G(s)$, we will use the strategy profile s and its corresponding created graph $G(s)$ interchangeably and we will say that $G(s)$ is in NE (or GE). Note that for any graph $G(s)$ in NE or GE, it follows that for every edge $(u, v) \in E_{G(s)}$ we have $(v, u) \notin E_{G(s)}$. This is true, since otherwise one of the agents could omit the other from its strategy without removing the edge from $G^\leftrightarrow(s)$ and thereby decrease its cost.

Given a temporal host graph H and a strategy profile s , we want to compare different created graphs in terms of their social cost. Here, the *social cost* of some created graph $G(s)$ is defined as

$$SC_H(s) = \sum_{v \in V_H} c_H(v, s) = |E_{G(s)}| + K \sum_{v \in V_H} |V_H \setminus R_{G^\leftrightarrow(s)}(v)|.$$

Note that the social cost of $G(s)$ equals $|E_{G(s)}|$ if $G^\leftrightarrow(s)$ is a temporal spanner. If for some host graph H the strategy profile s_H^* minimizes the social cost, we call s_H^* a *social optimum* for H and, by extension, the corresponding graph $G(s_H^*)$ a *social optimum subgraph* of H , often denoted as OPT. Since K is large, the set of social optimum subgraphs and the set of minimum temporal spanners for H coincide.

To investigate the efficiency loss from letting agents act selfishly towards minimizing their costs, we define the *Price of Anarchy (PoA)* [Koutsoupias and Papadimitriou, 1999]. The PoA is the worst ratio between the social cost of any stable state and the social cost of the corresponding social optimum on the same host graph. Towards a formal definition, let NE_H denote the set of strategy profiles that are in NE for a given host graph H . Moreover, let $\mathcal{H}_{n, t_{max}}$ denote the set of all possible complete temporal host graphs with n nodes and a lifetime of at most t_{max} . Then, the PoA, with respect to NE, is formally defined as

$$PoA_{NE}(n, t_{max}) = \sup_{H \in \mathcal{H}_{n, t_{max}}} \max_{s \in NE_H} \frac{SC_H(s)}{SC_H(s_H^*)}.$$

The PoA w.r.t. GE is defined analogously. We write $PoA_{NE}(n)$ instead of $PoA_{NE}(n, \binom{n}{2})$ when we consider host graphs with arbitrary lifetime.

Lastly, we define an important property of the game dynamics: We say that the TRNCG has the *finite improvement property* if any sequence of improving moves must be finite.

¹Formally, for $x \in S_v$ and $y \in V_H \setminus S_v$, we have $S'_v = S_v \setminus \{x\}$ or $S'_v = S_v \cup \{y\}$. Note, that typically in the literature, greedy improving moves also allow swaps, i.e. $S'_v = (S_v \setminus \{x\}) \cup \{y\}$. However, due to our model definition we can ignore swaps because if swapping x for y is improving, simply adding y is also improving.

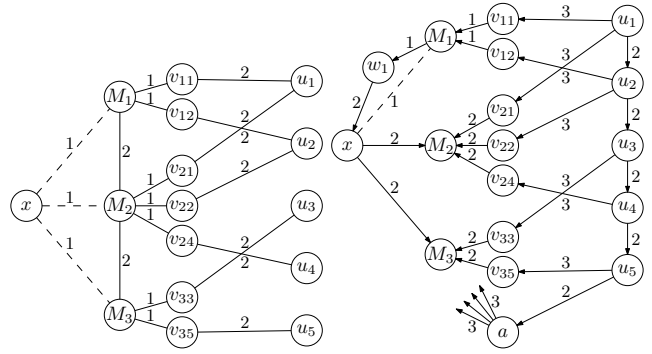


Figure 2: This figure shows examples of the constructions for Theorem 1 (left) and Theorem 2 (right), given the set cover instance consisting of the universe $U := \{u_1, \dots, u_5\}$ and the sets $\mathcal{M} := \{M_1, \dots, M_3\}$ with $M_1 := \{u_1, u_2\}$, $M_2 := \{u_1, u_2, u_4\}$ and $M_3 := \{u_3, u_5\}$. Additionally, on the right, we are given a set cover consisting of M_2 and M_3 but not M_1 . This is encoded by x buying edges towards M_2 and M_3 and the existence of w_1 .

2 Computational Complexity

In this section, we prove that computing best responses and checking strategy profiles whether they are in NE is NP-hard. This is a significant difference to network formation games with reachability objective where computing best responses is computationally easy [Bala and Goyal, 2000]. This means that adding the temporal component to the model makes it considerably harder.

Theorem 1. *Given a tuple (H, s, x) consisting of a complete temporal host graph H with lifetime $t \geq 2$, a strategy profile s , and a node $x \in V_H$, computing a best response for x is NP-hard.*

Theorem 2. *Deciding whether a pair (H, s) consisting of a complete temporal host graph H with lifetime $t \geq 3$ and a strategy profile s is a NE is NP-hard.*

Proof. Given a universe $U := \{u_1, \dots, u_k\}$ of k elements and a set of m sets $\mathcal{M} := \{M_1, \dots, M_m\} \subseteq \mathcal{P}(U)$, deciding whether a given set cover $\mathcal{C} \subseteq \mathcal{M}$ for U is a minimum set cover is NP-hard. We give a polynomial time reduction that, given an instance $(U, \mathcal{M}, \mathcal{C})$, constructs a tuple (H, s) consisting of a temporal host graph H and a strategy profile s . We show that \mathcal{C} is a minimum set cover if and only if s is a NE for H . Instead of defining s directly, we define $G := G(s)$. More precisely, the host graph H is defined as follows

$$V_H := \{x, a\} \cup \mathcal{M} \cup U \cup \bigcup_{i=1}^m \bigcup_{u_j \in M_i} \{v_{ij}\} \cup \bigcup_{M_i \in \mathcal{M} \setminus \mathcal{C}} \{w_i\}$$

$$\lambda(e) := \begin{cases} 1 & \text{if } \exists M_i \in \mathcal{M} \setminus \mathcal{C}, j \in \mathbb{N}: \\ & M_i \in e \wedge \{x, w_i, v_{ij}\} \cap e \neq \emptyset \\ 2 & \text{if } \exists M_i \in \mathcal{C}, j \in \mathbb{N}: \\ & (M_i \in e \wedge \{x, v_{ij}\} \cap e \neq \emptyset) \\ & \forall e \in \{\{w_j, x\}, \{u_n, a\}, \{u_j, u_{j+1}\}\} \\ 3 & \text{otherwise.} \end{cases}$$

The set of edges of $G(\mathbf{s})$ is defined as follows

$$\begin{aligned}
 E_G := & \bigcup_{i=1}^{n-1} \{(u_i, u_{i+1})\} \cup \bigcup_{i=1}^m \bigcup_{u_j \in M_i} \{(v_{ij}, M_i), (u_j, v_{ij})\} \\
 & \cup \{(u_n, a), (a, x)\} \cup \bigcup_{i=1}^m \{(a, M_i)\} \cup \bigcup_{M_i \notin \mathcal{C}} \{(a, w_i)\} \\
 & \cup \bigcup_{M_i \in \mathcal{C}} \{(x, M_i)\} \cup \bigcup_{M_i \notin \mathcal{C}} \{(M_i, w_i), (w_i, x)\}.
 \end{aligned}$$

Intuitively, we construct a node for each set in \mathcal{M} and each element in U and connect each set with all its elements via a monotonically increasing path of length 2. See Figure 2 for an example of the construction. The other edges are chosen so that G is a temporal spanner and all nodes except for x play best response. This can easily be checked for every node. Hence, to check whether \mathbf{s} is a NE, we only need to check whether x plays best response.

Let $\mathcal{C}' \subseteq \mathcal{M}$ be a minimum set cover for U and S_x^b a best response of x for \mathbf{s} .

Consider $S'_x := \mathcal{C}'$, meaning that x builds all the edges $\{x, M_i\}$ for $M_i \in \mathcal{C}'$. We see that x can now reach every node in $G(\mathbf{s}_{-x} \cup S'_x)$. Therefore, when \mathcal{C} is not a minimum set cover and therefore $|\mathcal{C}'| < |\mathcal{C}|$ it follows that $|S'_x| < |S_x|$ which implies that \mathbf{s} is not a NE.

Since S_x^b is a best response, x can reach every node in $G(\mathbf{s}_{-x} \cup S_x^b)$. Suppose, x buys an edge to one of the nodes u_j or v_{ij} . Instead, x can buy an edge to a node M_i , such that $u_j \in M_i$, without breaking reachability. Therefore, there is a best response $S_x^{b'} \subseteq \mathcal{M}$. Note that x still reaches all nodes $u \in U$ in $G(\mathbf{s}_{-x} \cup S_x^{b'})$ and u can only be reached by x if there is $M \in \mathcal{M}$ such that x builds an edge to M . This means that $\mathcal{C}'' := S_x^{b'}$ is a set cover. Therefore, if \mathbf{s} is not a NE, it follows $|S_x^{b'}| < |S_x|$ which implies that $|\mathcal{C}''| < |\mathcal{C}|$, so \mathcal{C} is not a minimum set cover.

It is obvious that this construction is computable in polynomial time which concludes the proof. \square

While NE is a very natural solution concept, the fact that it is computationally hard to compute best responses raises the question of whether it can realistically model the selfishness of the agents. Therefore, we will also consider GE because greedy best responses are computable in polynomial time.

Proposition 3. *Given a tuple (H, \mathbf{s}, x) consisting of a temporal host graph H , a strategy profile \mathbf{s} , and a node $x \in V_H$, a greedy best response for x is computable in polynomial time.*

3 Existence and Properties of Equilibria

We discuss under what circumstances equilibria exist and what properties they have. We start by showing that equilibrium existence cannot be proven via potential functions.

Theorem 4. *The TRNCG is not a potential game.*

Note that all the strategy changes in the improving response cycle are greedy best responses, too. This means that it is not a potential game even with regard to GE.

Next, we show that equilibria always exist when the lifetime is $t = 2$ and that we can find one in polynomial time.

This contrasts the result from Theorem 1 which showed that, even for $t = 2$, computing best responses is NP-hard.

Theorem 5. *Let H be a complete temporal host graph with $t = 2$. Then there is a strategy profile \mathbf{s} in NE for H .*

Proof. We show this by proving that H contains a spanning tree T whose edges all have the same label. Note that any strategy profile \mathbf{s} such that $G^{\leftrightarrow}(\mathbf{s}) = T$ is a NE.

Let H_i denote the subgraph of H on V_H that contains all edges of H of label i . We show that at least one between H_1 and H_2 is connected, thus proving the existence of T .

The claim trivially follows if H_1 is connected. So, assume that H_1 is not connected, i.e., there is a cut (X, Y) that is traversed by none of the edges in H_1 . All the edges that traverse the cut (X, Y) are in H_2 . Hence, H_2 is connected. \square

In the following, we prove that stable graphs cannot contain too many edges. We first bound the number of edges for graphs with a small lifetime t .

Theorem 6. *Let H be a complete temporal host graph containing $n > 2$ nodes and lifetime $t > 1$. Then any GE contains at most $t(n - 2)$ edges.*

Proof. Let \mathbf{s} be a strategy profile and $G := G(\mathbf{s})$ a GE. If there are at least n edges with some label l , some of them form a cycle. Removing one edge from the cycle does not change reachability among pairs of nodes. Therefore, each label can appear at most $n - 1$ times. Furthermore, if there are $n - 1$ edges and no cycles with the same label in G , those edges would form a spanning tree of G that guarantees reachability among pairs of nodes. Then, no other edge would be needed. Therefore, if G contains at least two labels, each label appears at most $(n - 2)$ times. Combined with only t labels existing, G contains at most $t(n - 2)$ edges. \square

Next, we prove an upper bound on the number of edges in an equilibrium state independent of t . We start by introducing the concept of *necessary edges* which we then use to characterize a structure that cannot appear in an equilibrium.

Definition 7 (necessary edge). *Let H be a complete temporal host graph with n agents, \mathbf{s} a strategy profile and $G := G(\mathbf{s})$. For each edge $e = (u, v) \in G$ that u buys, we define*

$$A_G(e) := \{x \in V \mid x \in R_G(u) \wedge x \notin R_{G-e}(u)\}.$$

We say that e is necessary for u to reach the agents in $A_G(e)$.

Note, that if G is a GE, $A_G(e) \neq \emptyset$ for all $e \in E_G$.

Using this definition, we characterize a structure that cannot appear in any strategy profile.

Lemma 8. *Let H be a complete temporal host graph, \mathbf{s} be a strategy profile, and $G := G(\mathbf{s})$. There cannot be nodes $z, u_1, u_2, x, y \in V$ and distinct edges $e_{1x}, e_{1y}, e_{2x}, e_{2y} \in E_G$ such that for $i \in \{1, 2\}$ and $j \in \{x, y\}$*

1. $\{z, u_i\} \in E_{G^{\leftrightarrow}} \setminus \{e_{ij}\}$;
2. e_{ij} is bought by u_i ;
3. $j \in A_G(e_{ij})$;
4. $\lambda(\{z, u_i\}) \leq \lambda(e_{ij})$.

(See Figure 3.)




Figure 3: A forbidden structure in a strategy profile. The node z has two neighbors u_1 and u_2 that both have two distinct edges that they need to reach the nodes x and y respectively. For both of them, the two needed edges have at least the same label as their edge to z .

Proof. Towards a contradiction, suppose there are nodes $z, u_1, u_2, x, y \in V$ and edges $e_{1x}, e_{1y}, e_{2x}, e_{2y} \in E_G$ as defined above. W.l.o.g. $\lambda(\{z, u_1\}) \leq \lambda(\{z, u_2\})$ and $\lambda(e_{1x}) \leq \lambda(e_{1y})$. Any temporal path from u_2 to x starts with e_{2x} (since $x \in A_G(e_{2x})$) and has to use e_{1x} . Otherwise, u_1 could reach x by using $\{z, u_1\}, \{z, u_2\}$ and then the path from u_2 to x without needing e_{1x} which contradicts $x \in A_G(e_{1x})$. The same holds for y instead of x .

Therefore, there is a temporal path P from u_2 to u_1 , which starts with e_{2x} and arrives at u_1 no later than $\lambda(e_{1x}) \leq \lambda(e_{1y})$. This means that u_2 does not need e_{2y} to reach y since it can use P to get to u_1 and travel to y from there. This contradicts $y \in A_G(e_{2y})$. \square

The forbidden structure from Figure 3 implies that graphs with at least $\sqrt{6n}^{\frac{3}{2}} + n$ edges must contain unnecessary edges, giving us a bound on the number of edges in an equilibrium.

Theorem 9. *Let H be a complete temporal host graph with $|V_H| = n$ agents and \mathbf{s} be a strategy profile. If $G := G(\mathbf{s})$ contains at least $\sqrt{6n}^{\frac{3}{2}} + n$ edges, then G is not a GE.*

Proof. Towards a contradiction, suppose that G is a GE and $|E_G| \geq \sqrt{6n}^{\frac{3}{2}} + n$. As shown in the full version [Bild et al., 2023], there is a node $z \in V_G$ and a set $M \subset V_G$ such that

1. $|M| = \lceil \frac{1}{3}\sqrt{6n} \rceil$;
2. $(u, z) \in E_G$, for every $u \in M$;
3. Each $u \in M$ has a set $E_u \subseteq E_G$ of at least $\frac{2}{3}\sqrt{6n}$ outgoing edges (u, v) with $z \neq v$ and $\lambda_G((u, z)) \leq \lambda((u, v))$.

See Figure 4 for an illustration.

For each edge $e \in E_u$, let $a_e \in A_G(e)$ be a representative of $A_G(e)$. Note that $A_G(e) \neq \emptyset$ because G is a GE, so those representatives always exist. For each $u \in M$, we define

$$D_u := \bigcup_{e \in E_u} \{a_e\}.$$




Figure 4: A structure that always appears in a directed temporal graph with at least $\sqrt{6n}^{\frac{3}{2}} + n$ edges. A node z exists with $\lceil \frac{\sqrt{6}}{3}\sqrt{n} \rceil$ neighbors via in-edges that each have at least $\frac{2\sqrt{6}}{3}\sqrt{n}$ out-edges with a label that is at least as high as the label of their edge to z .

Intuitively, D_u contains nodes that z can reach by going over u and that u needs to buy an edge for.

We see that the forbidden structure from Lemma 8 appears if there are two nodes $u, v \in M$ such that $|D_u \cap D_v| \geq 2$. We can therefore assume $|D_u \cap D_v| \leq 1$ for all $u, v \in M$. Also, for $e, e' \in E_u$ we have $A_G(e) \cap A_G(e') = \emptyset$ since there cannot be two edges that are necessary for u to reach the same node. From this, we get $|D_u| \geq |E_u| \geq \frac{2}{3}\sqrt{6n}$.

Using the inclusion-exclusion principle, we get

$$\begin{aligned} \left| \bigcup_{u \in M} D_u \right| &\geq \sum_{u \in M} |D_u| - \sum_{\{u, v\} \subseteq M, u \neq v} |D_u \cap D_v| \\ &\geq \left\lceil \frac{1}{3}\sqrt{6n} \right\rceil \frac{2}{3}\sqrt{6n} - \frac{1}{2} \left\lceil \frac{1}{3}\sqrt{6n} \right\rceil \left(\left\lceil \frac{1}{3}\sqrt{6n} \right\rceil - 1 \right) \\ &> \left\lceil \frac{1}{3}\sqrt{6n} \right\rceil \frac{2}{3}\sqrt{6n} - \frac{1}{2} \left\lceil \frac{1}{3}\sqrt{6n} \right\rceil \frac{1}{3}\sqrt{6n} \\ &= \left\lceil \frac{1}{3}\sqrt{6n} \right\rceil \frac{1}{2}\sqrt{6n} \geq n. \end{aligned}$$

This is a contradiction since G has only n nodes. \square

4 Quality of Equilibria

Finally, we characterize the quality of equilibria by analyzing the Price of Anarchy (PoA). We show that the PoA with regard to NE and GE are within a $\log(n)$ -factor of each other, meaning that any bound on the easier to analyze Greedy Equilibria also yields a bound for Nash Equilibria. We proceed to give several upper and lower bounds.

For analyzing the PoA, we start by upper bounding the cost of the social optimum. The following result follows from [Casteigts et al., 2021].

Theorem 10. *Let H be a complete temporal host graph with $|V_H| = n$ agents and OPT be a social optimum for H . Then $\text{SC}_H(\text{OPT}) = |E_{\text{OPT}}| \in \mathcal{O}(n \log(n))$.*

Whether we consider NE or GE has an impact on the PoA. However, we show that these two PoA's differ only by at most a factor of $O(\log(n))$.

Theorem 11. $\text{PoA}_{\text{GE}}(n) \leq \mathcal{O}(\log(n)) \text{PoA}_{\text{NE}}(n)$.

Proof. Let H be a complete temporal host graph with n vertices and lifetime t . Moreover, let \mathbf{s} be a strategy profile in GE and let \mathbf{s}_H^* the social optimum for H such that $\text{PoA}_{\text{GE}}(n) = \frac{\text{SC}_H(\mathbf{s})}{\text{SC}_H(\mathbf{s}_H^*)}$. We construct a new temporal host graph H' by relabeling all edges not in $G(\mathbf{s})$ to $t + 1$.

First, we argue that \mathbf{s} is in NE with respect to H' . Consider a strategy change for a node u that removes k edges and adds l edges. Since \mathbf{s} is stable against edge removal, removing those k edges will lead to at least k nodes no longer being reachable from u . Additionally, agent u can use the l added edges only to reach the other endpoints of the edges since all edges not in $G(\mathbf{s})$ have a higher label in H' than all the edges in $G(\mathbf{s})$. Therefore, for agent u to reach everyone after the strategy change, we need $k \leq l$ which means that the strategy change is not an improving move.

We also have $\text{SC}_H(\mathbf{s}) = \text{SC}_{H'}(\mathbf{s})$ and $\text{SC}_H(\mathbf{s}_H^*) \geq n - 1$. Let $\mathbf{s}_{H'}^*$ be a social optimum for H' . By Theorem 10

we have $SC_{H'}(\mathbf{s}_{H'}^*) \in \mathcal{O}(n \log(n))$, from which we get $SC_H(\mathbf{s}_H^*) \in \Omega\left(\frac{SC_{H'}(\mathbf{s}_{H'}^*)}{\log(n)}\right)$ and therefore

$$\frac{SC_H(\mathbf{s})}{SC_H(\mathbf{s}_H^*)} \leq \frac{\mathcal{O}(\log(n)) SC_{H'}(\mathbf{s})}{SC_{H'}(\mathbf{s}_{H'}^*)} \leq \mathcal{O}(\log(n)) \text{PoA}_{\text{NE}}(n)$$

which proves the claim. \square

When considering the PoA with respect to a fixed maximum lifetime t_{max} instead of a fixed n , we get an upper bound for the PoA of t_{max} .

Theorem 12. *For fixed maximum lifetime t_{max} , we have $\text{PoA}_{\text{GE}}(n, t_{max}) \leq t_{max} - \frac{t_{max}}{n-1}$. For $t_{max} = 2$, we have $\text{PoA}_{\text{GE}}(n, 2) = \text{PoA}_{\text{NE}}(n, 2) = 2 - \frac{2}{n-1}$.*

Proof sketch. The upper bound on $\text{PoA}_{\text{GE}}(n, t_{max})$ follows directly from Theorem 6 by lower bounding the social cost of the optimum with $n - 1$.

For the tight lower bound when $t_{max} = 2$ which also holds for $\text{PoA}_{\text{NE}}(n, 2)$, we construct the graph from Figure 5. \square

We show that the PoA scales at least logarithmically with n by giving a class of host graphs that have equilibria which are by a logarithmic factor worse than the social optimum.

Theorem 13. *It holds that $\text{PoA}_{\text{NE}}(n) \in \Omega(\log n)$.*

Proof. Let $n \geq 8$ be a power of 2. Kempe *et al.* [2002] proved how to label a $\log(n)$ -dimensional hypercube with labels in $\{1, \dots, \log(n)\}$ so as the resulting temporal graph is a minimal temporal spanner. The idea of our proof is to define the host graph H on n nodes so that it contains the $\log(n)$ -dimensional hypercube which is a NE regardless of the strategy profile that generates it, and a spanning tree formed by edges having the same label 1 + $\log(n)$.

Formally, the nodes are bitstrings of length $\log(n)$. The complete host graph H is defined as follows

$$V_H := \{b_1 b_2 \dots b_{\log(n)} \mid \forall 1 \leq i \leq \log(n): b_i \in \{0, 1\}\}$$

$$\lambda(\{u, v\}) := \begin{cases} i & \text{if } u \text{ and } v \text{ differ only in position } i; \\ \log(n) + 1 & \text{otherwise.} \end{cases}$$

The strategy profile \mathbf{s} induces a graph $G := G^{\leftrightarrow}(\mathbf{s})$ with

$$E_{G^{\leftrightarrow}} := \{\{u, v\} \mid v \text{ and } u \text{ differ in exactly one bit}\}.$$




Figure 5: Left: A NE containing $2(n - 2)$ edges and only 2 labels. All edges not taken have label 2. Right: A 3-dimensional hypercube. Each node corresponds to a bitstring of length 3. An edge exists between two bitstrings if they differ in exactly one position. The label of an edge is the position that the incident bitstrings differ in.

Figure 5 shows an illustration of G for $n = 8$.

Since $n \geq 8$, H contains a spanning tree OPT whose edges are all labelled with $\log(n) + 1$. Since OPT is a temporal spanner, it is a social optimum with $SC_H(\text{OPT}) = n - 1$.

G is a hypercube graph containing $\frac{n}{2} \log(n)$ edges and therefore $SC_H(G) = \frac{n}{2} \log(n)$. If G is a NE, we get the desired bound

$$\text{PoA}_{\text{NE}}(n) \geq \frac{SC_H(G)}{SC_H(\text{OPT})} = \frac{\frac{n}{2} \log(n)}{n - 1} \in \Omega(\log n).$$

Finally, we argue that G is a NE. First, we observe that there is exactly one temporal path from every node u to every other node v . When considering the bit strings of these two nodes, the path flips all of the bits that are different in u and v in ascending order. By definition of H , all of those edges exist and are labeled in ascending order.

Secondly, note that all edges $\{u, v\} \in E_{G^{\leftrightarrow}}$ are needed by both of their end points to reach each other. By definition of E_G , the bitstrings of u and v differ only in one position p . When removing the edge $\{u, v\}$, a temporal path starting from u has to flip another bit than p . In a temporal path on G , this bit can never be flipped back, so v cannot be reached. We also observe that buying other edges outside of E_G cannot replace any edge in E_G .

Hence, G is a temporal spanner and no agent can improve their strategy, which makes \mathbf{s} a NE. \square

An upper bound for the PoA follows from Theorem 9.

Corollary 14. $\text{PoA}_{\text{GE}}(n) \in \mathcal{O}(\sqrt{n})$.

Proof. This follows from Theorem 9 and by lower bounding the social cost of the optimum with $n - 1$. \square

5 Conclusion and Outlook

In this paper, we combine game-theoretic network creation with temporal graphs. To this end, we defined and analyzed the Temporal Reachability Network Creation Game.

Even though we consider a restricted setting with unit cost on each edge and a complete host graph, we show NP-hardness for computing best responses and for deciding NE, showing that adding temporal aspects to the model makes it much harder. As our main contribution, we show non-trivial structural properties of equilibria and use them to derive several upper and lower bounds on the Price of Anarchy.

Since the upper bound of $\mathcal{O}(\sqrt{n})$ on the PoA only uses one local property, we believe that the PoA is closer to our lower bound of $\Omega(\log(n))$. Another important open question is settling the existence of equilibria for all complete temporal host graphs. We conjecture that equilibria exist, but, based on our efforts, even proving this for lifetime $t = 3$ is challenging.

We laid the groundwork for future research in this field. There are many natural extensions of our model. As far as agent strategy is concerned, the agents might want to minimize the distance to all others or, due to the time attribute introduced by our model, the agents may want to minimize their arrival time at the other agents. Also structural properties of the host graph could be altered. For enhanced realism, the edges could be directed, have non-uniform buying costs, and/or non-instant traversal times. The rules of the game can also be adjusted, for example by allowing cooperation.

References

- [Ahmed *et al.*, 2020] Abu Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen G. Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Comput. Sci. Rev.*, 37:100–253, 2020.
- [Andelman *et al.*, 2009] Nir Andelman, Michal Feldman, and Yishay Mansour. Strong price of anarchy. *Games Econ. Behav.*, 65(2):289–317, 2009.
- [Axiotis and Fotakis, 2016] Kyriakos Axiotis and Dimitris Fotakis. On the size and the approximability of minimum temporally connected subgraphs. In *ICALP 2016*, pages 149:1–149:14, 2016.
- [Bala and Goyal, 2000] Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.
- [Bilò *et al.*, 2016] Davide Bilò, Luciano Gualà, Stefano Leucci, and Guido Proietti. Locality-based network creation games. *ACM Trans. Parallel Comput.*, 3(1):6:1–6:26, 2016.
- [Bilò *et al.*, 2019] Davide Bilò, Tobias Friedrich, Pascal Lenzner, and Anna Melnichenko. Geometric network creation games. In *SPAA 2019*, pages 323–332, 2019.
- [Bilò *et al.*, 2021] Davide Bilò, Tobias Friedrich, Pascal Lenzner, Stefanie Lowski, and Anna Melnichenko. Selfish creation of social networks. In *AAAI 2021*, pages 5185–5193, 2021.
- [Bilò *et al.*, 2022] Davide Bilò, Gianlorenzo D’Angelo, Luciano Gualà, Stefano Leucci, and Mirko Rossi. Sparse temporal spanners with low stretch. In *ESA 2022*, pages 19:1–19:16, 2022.
- [Bilò *et al.*, 2023] Davide Bilò, Sarel Cohen, Tobias Friedrich, Hans Gawendowicz, Nicolas Klodt, Pascal Lenzner, and George Skretas. Temporal network creation games. Technical Report 2305.07494, arXiv, 2023. Full version of this paper.
- [Bullinger *et al.*, 2022] Martin Bullinger, Pascal Lenzner, and Anna Melnichenko. Network creation with homophilic agents. In *IJCAI 2022*, pages 151–157, 2022.
- [Casteigts *et al.*, 2021] Arnaud Casteigts, Joseph G. Peters, and Jason Schoeters. Temporal cliques admit sparse spanners. *J. Comput. Syst. Sci.*, 121:1–17, 2021.
- [Chauhan *et al.*, 2016] Ankit Chauhan, Pascal Lenzner, Anna Melnichenko, and Martin Münn. On selfish creation of robust networks. In *SAGT 2016*, pages 141–152, 2016.
- [Chauhan *et al.*, 2017] Ankit Chauhan, Pascal Lenzner, Anna Melnichenko, and Louise Molitor. Selfish network creation with non-uniform edge cost. In *SAGT 2017*, pages 160–172, 2017.
- [Chen *et al.*, 2019] Yu Chen, Shahin Jabbari, Michael J. Kearns, Sanjeev Khanna, and Jamie Morgenstern. Network formation under random attack and probabilistic spread. In *IJCAI 2019*, pages 180–186, 2019.
- [Corbo and Parkes, 2005] Jacomo Corbo and David C. Parkes. The price of selfish behavior in bilateral network formation. In *PODC 2005*, pages 99–107, 2005.
- [Cord-Landwehr and Lenzner, 2015] Andreas Cord-Landwehr and Pascal Lenzner. Network creation games: Think global - act local. In *MFCS 2015*, pages 248–260, 2015.
- [Echzell *et al.*, 2020] Hagen Echzell, Tobias Friedrich, Pascal Lenzner, and Anna Melnichenko. Flow-based network creation games. In *IJCAI 2020*, pages 139–145, 2020.
- [Eidenbenz *et al.*, 2006] Stephan J. Eidenbenz, V. S. Anil Kumar, and Sibylle Züst. Equilibria in topology control games for ad hoc networks. *Mob. Networks Appl.*, 11(2):143–159, 2006.
- [Fabrikant *et al.*, 2003] Alex Fabrikant, Ankur Luthra, Elitza N. Maneva, Christos H. Papadimitriou, and Scott Shenker. On a network creation game. In *PODC 2003*, pages 347–351. ACM, 2003.
- [Friedrich *et al.*, 2017] Tobias Friedrich, Sven Ihde, Christoph Keßler, Pascal Lenzner, Stefan Neubert, and David Schumann. Efficient best response computation for strategic network formation under attack. In *SAGT 2017*, pages 199–211, 2017.
- [Friedrich *et al.*, 2022] Tobias Friedrich, Hans Gawendowicz, Pascal Lenzner, and Anna Melnichenko. Social distancing network creation. In *ICALP 2022*, pages 62:1–62:21, 2022.
- [Goyal *et al.*, 2016] Sanjeev Goyal, Shahin Jabbari, Michael J. Kearns, Sanjeev Khanna, and Jamie Morgenstern. Strategic network formation with attack and immunization. In *WINE 2016*, pages 429–443, 2016.
- [Gulyás *et al.*, 2015] András Gulyás, József J Bíró, Attila Kőrösi, Gábor Rétvári, and Dmitri Krioukov. Navigable networks as Nash equilibria of navigation games. *Nature communications*, 6(1):1–10, 2015.
- [Jackson and Wolinsky, 1996] Matthew O. Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996.
- [Kempe *et al.*, 2002] David Kempe, Jon Kleinberg, and Amit Kumar. Connectivity and inference problems for temporal networks. *Journal of Computer and System Sciences*, 64(4):820–842, 2002.
- [Klobas *et al.*, 2022] Nina Klobas, George B. Mertzios, Hendrik Molter, and Paul G. Spirakis. The complexity of computing optimum labelings for temporal connectivity. In *MFCS 2022*, pages 62:1–62:15, 2022.
- [Koutsoupias and Papadimitriou, 1999] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *STACS 1999*, pages 404–413, 1999.
- [Lenzner, 2012] Pascal Lenzner. Greedy selfish network creation. In *WINE 2012*, pages 142–155, 2012.
- [Meirom *et al.*, 2015] Eli A. Meirom, Shie Mannor, and Ariel Orda. Formation games of reliable networks. In *INFOCOM 2015*, pages 1760–1768, 2015.

[Narasimhan and Smid, 2007] Giri Narasimhan and Michiel Smid. *Geometric spanner networks*. Cambridge University Press, 2007.

[Nash, 1950] John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.