Econometrica, Vol. 87, No. 2 (March, 2019), 567–591

PEER-CONFIRMING EQUILIBRIUM

ELLIOT LIPNOWSKI
Department of Economics, University of Chicago

EVAN SADLER
Department of Economics, Columbia University

We can often predict the behavior of those closest to us more accurately than that of complete strangers, yet we routinely engage in strategic situations with both: our social network impacts our strategic knowledge. Peer-confirming equilibrium describes the behavioral consequences of this intuition in a noncooperative game. We augment a game with a network to represent strategic information: if two players are linked in the network, they have correct conjectures about each others' strategies. In peer-confirming equilibrium, there is common belief that players (i) behave rationally and (ii) correctly anticipate neighbors' play. In simultaneous-move games, adding links to the network always restricts the set of outcomes. In dynamic games, the outcome set may vary non-monotonically with the network because the actions of well-connected players help poorly-connected players coordinate. This solution concept provides a useful language for studying public good provision, highlights a new channel through which central individuals facilitate coordination, and delineates possible sources of miscoordination in protests and coups.

KEYWORDS: Networks, strategic uncertainty, conjectural equilibrium, forward induction.

1. INTRODUCTION

SOCIAL NETWORKS are important in a wide range of economic contexts. From peer effects to innovation adoption to job searches, our social ties—both strong and weak—have a profound impact on what we know and how we act.[1] These relationships may also affect our expectations of others in strategic interactions. Our everyday experiences guide our expectations for the future, and we inevitably interact with some people far more frequently than with others. This suggests that our conjectures about close friends and neighbors should be more accurate than those about complete strangers. Social networks therefore have a natural role in describing how individuals coordinate their actions. Our paper explores the relationship between social network structure and strategic coordination.

We introduce a solution concept, *peer-confirming equilibrium*, that makes the dependence of strategic information on social ties explicit. A network, represented as an undirected graph G, describes players' strategic information and is part of the basic data of the game. Players respond optimally to conjectures about other players' strategies, and a link between two players indicates that their conjectures about each other are accurate.

Elliot Lipnowski: lipnowski@uchicago.edu
Evan Sadler: es3668@columbia.edu

We are indebted to Dan Bernhardt, Emiliano Catonini, Arun Chandrasekhar, Amanda Friedenberg, Annie Liang, Phil Reny, Richard van Weelden, and Wenjun Wang for useful comments and suggestions. We also thank Ben Brooks, Tommaso Denti, Ben Golub, Hanna Halaburda, Aditya Kuvalekar, Doron Ravid, and participants at the third annual Conference on Network Science in Economics and the 2017 European Winter Meeting of the Econometric Society for their comments. We thank Eric Karsten for research assistance.

[1]See Jackson (2008) and Bramoullé, Galeotti, and Rogers (2016) for recent surveys of the literature on the economics of networks.

 https://doi.org/10.3982/ECTA15597

Unlike much of the economics literature on networks, the network *does not* indicate any relationship between players' payoff functions nor a channel for communication. A strategy profile constitutes a peer-confirming equilibrium for the network if it is consistent with three criteria:

(a) Each player responds optimally to her conjecture about others' strategies.
(b) Players' conjectures about neighbors' behavior are correct.
(c) The above two facts are common belief among the players.

Strategic uncertainty is greatest when the network is empty: no one knows anyone else's strategy. In this case, our definition yields exactly the set of rationalizable strategy profiles. The outcome set in a simultaneous-move game is monotone in the network: adding links restricts the set of permissible profiles. At the other extreme, a complete graph, our solution concept is equivalent to Nash equilibrium.

Peer-confirming equilibrium highlights a new relationship between social networks and strategic behavior. In several examples, we explore how the network structure and other details of the game jointly influence the set of equilibria. For instance, the action of a fully connected player can convey information about others' choices, but with different payoff structures this may or may not lead to coordination on a Nash equilibrium. In a public goods provision game, with each player simultaneously making a contribution, a single player with links to all others can induce Nash equilibrium play. If the fully connected player contributes a positive amount, her choice implicitly communicates the true marginal value of contributions to all other players. This contrasts sharply with an example we call "the potluck game," in which players try to bring as many different dishes as possible to a party, and a fully connected player's action conveys relatively little information. A fully connected player does more to aid coordination when the action space is rich and best responses are appropriately sensitive to others' behavior.

We then consider an application to political protests or coups, modeled as binary action coordination games (e.g., Angeletos, Hellwig, and Pavan (2007)). The success of a protest depends upon the effective coordination of many dispersed agents, succeeding only if enough of them join the effort. If not, those who participate may pay a high price for opposing the entrenched regime. Coordination is clearly a challenge, as history is replete with examples of failed coups and repressed political movements.[2] Instead of studying uncertainty about payoffs or about the exogenous "strength" of the regime, we consider uncertainty about others' strategies. In this case, fully connected players are key: a single player linked to all others leads to Nash equilibrium play. However, this coordination is fragile. If we remove a single link from the fully connected player, failed coups can occur in a peer-confirming equilibrium. Nevertheless, the network structure can still offer insight into the behavioral patterns we might expect. In a peer-confirming equilibrium, any two players whose neighborhoods jointly cover the entire population must coordinate on the same action. This suggests that social "elites" will coordinate with each other even if the population as a whole fails to do so.

Following these examples, we develop the analogous solution concept for dynamic games. Our definition coincides with subgame perfect equilibrium in a complete network and extensive-form rationalizability in an empty network. Neither of these concepts nests the other, so for dynamic games the outcome set is no longer monotone in the network. This reflects a subtle interaction between the network structure and forward induction reasoning. A player can use her action to signal the intentions of *other players* to whom

[2]Powell and Thyne (2011) comprehensively document hundreds of coups and coup attempts around the world over the last several decades, roughly half of which were successful.

she is connected: assuming she made a rational choice, her opponents make inferences about her strategic information. If the other players are not themselves well-connected, they have no information to contradict this signal, and our player can credibly convey information about what others will do. Adding links to the network may actually expand the set of peer-confirming equilibria because it inhibits this kind of strategic signaling.

Signaling of strategic information allows us to refine the predictions of both subgame perfect equilibrium and extensive-form rationalizability in a two-stage version of the protest game. Suppose a "leader" first publicly commits to protest or not, and only then do the remaining players choose whether to join. In an empty network, every profile is rationalizable. In a complete network, there are multiple equilibria—with a successful coup in some but not others. In contrast, a star network centered on the leader permits only one peer-confirming equilibrium outcome: all players participate in a successful protest. In this network, the leader's choice to protest is a convincing signal that the protest will succeed. Knowing this, the leader always protests because doing so induces others to join.

Beyond our conceptual contribution, peer-confirming equilibrium offers an empirically relevant tool to refine rationalizability—one with clear testable implications. For instance, to give subjects appropriate feedback in repeated play of a simultaneous-move game, an experimenter could provide information on the actions of particular individuals, according to a chosen network structure. Our theory predicts a smaller range of outcomes when the network is more dense, and by varying the payoff structure in a coordination game, one could test whether the presence of a fully connected player facilitates coordination as predicted. Outside the laboratory, social ties give us a meaningful and systematic way to restrict conjectures. From sharing gossip to sending marketing referrals to coordinating political activism, knowledge of friends' typical behavior is an important input for our own decisions. Increasingly, economists are able to measure social networks, and our solution concept suggests new avenues to make use of these data. Through our framework, one can assess the players and connections that are most crucial for coordination and use this to inform policy.

1.1. *Related Work*

Our paper introduces a new dimension to the growing literature on economics and social networks. Empirically, social ties are an important source of information (Marmaros and Sacerdote (2002), Banerjee, Chandrasekhar, Duflo, and Jackson (2013)), and the choices of friends and neighbors often have spillover effects (Bandiera and Rasul (2006), Carrell, Sacerdote, and West (2013)). Inspired by these findings, theoretical work extensively studies network games with externalities (Ballester, Calvó-Armengol, and Zenou (2006), Bramoullé and Kranton (2007), Galeotti, Goyal, Jackson, Vega-Redondo, and Yariv (2010)), social learning in networks (Golub and Jackson (2010), Lobel and Sadler (2015)), and models of network formation (Jackson and Wolinsky (1996), Galeotti and Goyal (2010)). We explore the implications of social structure for strategic coordination. Unlike the literature on network games, our network does not represent relationships between players' payoff functions, and unlike the literature on social learning, our network does not represent conduits for information flows. The connections we study represent knowledge about how other people intend to act. Actions can and do convey information to others, particularly in dynamic games, creating a superficial similarity to observational learning models. However, the transmission of information implied in peer-confirming equilibrium is less direct than this: there need not be any private signals about which players make inferences. Rather, the relevant information is implicit in players' conjectures about each other.

Our solution concept augments a game with a network to describe players' information concerning each other's actions. This contrasts with related work by Chwe (2000), in which a network describes players' beliefs concerning others' payoff types. It also contrasts with work by Tsakas (2013) and Bach and Tsakas (2014), which—providing foundations for existing solution concepts—employ a network to describe players' conjectures concerning each other's beliefs. While the goal in their work is quite different from ours, a common theme emerges: local conditions on players' strategic knowledge can have global consequences.

Through our applications, we also relate to a large literature that studies the coordination of political protests and coups. Much of this work applies the framework of global games (Carlsson and van Damme (1993), Morris and Shin (1998)), where the information contained in private signals about regime strength determines the extent of coordination. Recent contributions show how dynamic learning can lead to a multiplicity of equilibria, with alternating periods of peace and protest activity (Angeletos, Hellwig, and Pavan (2007)), and how an entrenched regime may use propoganda to its advantage, even when private information is precise (Edmond (2013)). Rather than focus on uncertainty about fundamental information (e.g., information about regime strength), we isolate the role of strategic uncertainty. Both likely play an important role in political movements, and both are therefore important to our understanding of these phenomena.

Formally, our solution concept is a specialization of rationalizable conjectural equilibrium (Rubinstein and Wolinsky (1994)). Rubinstein and Wolinsky motivated RCE as the result of learning through repeated play when individuals receive only coarse signals about the outcome in each iteration. As Esponda (2013) describes, RCE "intends to capture the steady state of a learning process that combines learning and introspection." In a peer-confirming equilibrium, these signals correspond to the strategies of a subset of players: we can interpret peer-confirming equilibrium as what should result when we limit feedback to observing neighbors' behavior. Our conceptual contribution is to link individuals' feedback to an underlying social network, suggesting a foundation for players' expectations about one another. More recently, Fudenberg and Kamada (2015) introduce another related concept, rationalizable partition-confirmed equilibrium. In simultaneous-move games, this concept is equivalent to rationalizable conjectural equilibrium, and so our concept specializes this as well. However, in dynamic games, our players engage in forward induction reasoning, so the concepts are distinct. Peer-confirming equilibrium provides a natural lens through which to examine situations in which most strategic information comes from our peers. While some earlier work focuses on the robustness of equilibrium to different forms of feedback, asking what information is required to reach Nash equilibrium, ours focuses on alternative predictions that peer-confirming equilibrium can furnish.

Work on rationalizability (Pearce (1984), Bernheim (1984), Battigalli and Siniscalchi (2002)) is a clear antecedent to our study of sophisticated play under strategic uncertainty, as is the broader literature on epistemic game theory (Dekel and Siniscalchi (2015)). Some of this research explores the consequences of rational play together with various belief restrictions (e.g., Aumann and Brandenburger (1995), Battigalli and Siniscalchi (2003)), while other work examines joint restrictions on players' beliefs and realized play (e.g., Section 5 of Battigalli and Siniscalchi (2002), Friedenberg (2017), Catonini (2017)). Although we do not conduct a formal epistemic analysis, our paper is closer to the latter stream. An important distinction between the epistemic approach and ours is that our "additional parameter" is a much simpler object: the collection of all strategic type spaces for a game is infinite dimensional and varies with the game form, while the collection of networks is finite and varies only with the player set.

2. PCE IN SIMULTANEOUS-MOVE GAMES

Let $\Gamma = \langle N, (S_i)_{i\in\mathbb{N}}, (u_i)_{i\in\mathbb{N}}\rangle$ denote a simultaneous-move game with complete information. The set of players is N, the set of strategies for player i is S_i, and the payoff function for player i is u_i.[3] In all of our applications, we take S_i to be player i's set of pure strategies.[4] We augment Γ with a network structure, which we represent as an undirected graph G. We write $ij \in G$ if players i and j are linked, and we write G_i for the set of all other players to whom i is linked. We impose strategic sophistication and use the graph G to further restrict the set of permissible strategy profiles. This section formalizes a solution concept in which players best respond to conjectures about others' strategies, players have correct conjectures about their neighbors in G, and these two facts are commonly believed.

To properly state our definition, we introduce notation to talk about players' conjectures and best replies. Let $S = \prod_{i\in N} S_i$ denote the set of strategy profiles, and let $S_{-i} = \prod_{j\neq i} S_j$ denote the set of partial profiles of strategies for players other than i. A **conjecture** for player $i \in N$ is a probability distribution $\mu_i \in \Delta(S_{-i})$, representing beliefs about what the other players will do. A strategy $s_i^* \in S_i$ is a **best reply** to the conjecture μ_i if

$$s_i^* \in \arg\max_{s_i\in S_i} \int_{S_{-i}} u_i(s_i, \cdot)\,\mathrm{d}\mu_i.$$

We write $r_i(\mu_i) \subseteq S_i$ for the set of all best replies to μ_i; this is the set of strategies player i can rationally adopt if her beliefs about other players' strategies are given by μ_i.

The graph G places restrictions on players' conjectures, and we need additional notation to state these restrictions. For each player $i \in N$ and strategy profile $\sigma \in S$, we define the set of **strategies consistent with i's strategic information at σ** as

$$S_{-i}^{\sigma,G} = \{s_{-i} \in S_{-i} : s_j = \sigma_j \ \forall j \in G_i\}.$$

In words, the set $S_{-i}^{\sigma,G}$ is the set of partial profiles in which i's neighbors in G take actions consistent with σ. When σ is played, the set $S_{-i}^{\sigma,G}$ represents the set of profiles to which i might assign positive probability in her conjecture μ_i: beliefs about neighbors' actions are correct. Similarly, for $\sigma \in \Sigma \subseteq S$, we define

$$\Sigma_{-i}^{\sigma,G} = S_{-i}^{\sigma,G} \cap \bigl\{s_{-i} \in S_{-i} : (\sigma_i, s_{-i}) \in \Sigma\bigr\}.$$

The set $\Sigma_{-i}^{\sigma,G}$ is the subset of $S_{-i}^{\sigma,G}$ that is consistent with the collection of strategy profiles Σ. If i is sure that the true strategy profile is in Σ, and σ is the true profile, then $\Sigma_{-i}^{\sigma,G}$ is the set of profiles to which i can assign positive probability in her conjecture μ_i.

Finally, we use the above notation to describe a best reply map for the network G. For each player $i \in N$, each measurable $\Sigma \subseteq S$, and each $\sigma \in \Sigma$, we define the set of player i's **viable conjectures relative to Σ at σ** as

$$\Delta_i^{\sigma,G}(\Sigma) = \bigl\{\mu_i \in \Delta(S_{-i}) : \mu_i\bigl(\Sigma_{-i}^{\sigma,G}\bigr) = 1\bigr\}.$$

[3]Here, N is finite, each S_i is a measurable space, and each u_i is a bounded, measurable function.

[4]A user of our solution concept may wish to study a mixed extension of a game, letting $S_i = \Delta A_i$ be a set of mixed strategies. In this case, our definition implicitly assumes that all players randomize independently. In particular, i having correct beliefs about neighbor j's play (as peer-confirming equilibrium requires) does more than impose correct marginal beliefs over a_j, but also implies that i views a_j as independent of others' play. We thank an anonymous referee for this observation.

This is the set of conjectures for player i that place probability 1 on a strategy profile in Σ and on i's neighbors playing σ.[5] The set of **network-consistent best replies** to Σ is

$$B_G(\Sigma) = \{\sigma \in \Sigma : \forall i \in N, \exists \mu_i \in \Delta_i^{\sigma,G}(\Sigma) \text{ s.t. } \sigma_i \in r_i(\mu_i)\}.$$

An element of $B_G(\Sigma)$ is a strategy profile in which each player best responds to a conjecture that is consistent with her neighbors playing σ and everyone playing something in Σ. The function $B_G(\Sigma)$ acts as a best reply correspondence on the space of strategy profiles, and with it we can now define the set of **peer-confirming equilibria**.

DEFINITION 1: Given a simultaneous-move game Γ and a network G, a strategy profile σ is a **(rationalizable) peer-confirming equilibrium** if it is contained in a measurable set $\Sigma \subseteq S$ such that $\Sigma = B_G(\Sigma)$. We write R_G for the set of all peer-confirming equilibria for the network G.

In a peer-confirming equilibrium, every player best responds to a conjecture that is correct about her neighbors' strategies and consistent with all others playing a peer-confirming equilibrium. Observe that the map B_G is monotone with respect to set inclusion: if Σ is smaller, the set of viable conjectures is smaller, so that the set of best replies is smaller. For sufficiently well-behaved games, this means we can compute R_G through a process of iterated deletion of strategy profiles. For instance, if S_i is compact and u_i is continuous for each i, then an equivalent definition of R_G is[6,7]

$$R_G \equiv \bigcap_{k=0}^{\infty} B_G^k(S).$$

A few properties are immediate from the definition.

PROPOSITION 1: *In a simultaneous-move game, the set of peer-confirming equilibria R_G satisfies the following*:

(a) *If G is empty, then R_G is the set of (correlated) rationalizable strategy profiles.*
(b) *If G is complete, then R_G is the set of Nash equilibria.*
(c) *If $G \subseteq G'$, then $R_G \supseteq R_{G'}$.*

As we add links to the network G, the set of conjectures that players can entertain shrinks. With fewer permissible conjectures, there are fewer best replies. In an empty graph, the definition is formally equivalent to the set of (correlated) rationalizable strategies (Pearce (1984), Bernheim (1984)): players can form any conjecture that is consistent with common belief of rationality. At the other extreme, a complete network implies that all players best respond to the actual strategies being played, and we reach Nash equilibrium. In between, the structure of the network creates nontrivial restrictions on the set of permissible strategy profiles.

[5]Notice that players are allowed to hold correlated conjectures, that is, elements of $\Delta_i^{\sigma,G}(\Sigma)$ need not be product measures. Just as is the case for correlated rationalizability, some actions of player i may be a best response *only* to such correlated conjectures.

[6]Here, B_G^k denotes the composition map. So $B_G^0(S) = S$, and $B_G^k(S) = B_G(B_G^{k-1}(S))$ for every $k \in \mathbb{N}$.

[7]The equivalence follows from monotonicity of B_G and upper hemicontinuity of the best response correspondences $\{r_i\}_{i\in N}$.

A network of relationships can serve more than one role. There is typically significant overlap between individuals' economic relationships, as in the literature on network games (e.g., Galeotti et al. (2010)), and their social ties. We say j is **payoff-irrelevant** to i if $u_i(\sigma_j, \sigma_{-j}) = u_i(s_j, \sigma_{-j})$ for all strategy profiles σ and all actions $s_j \in S_j$. Otherwise, player j is **payoff-relevant** to i. The network of payoff-relevance $\tilde{G}$—with an edge between players if and only if one is payoff-relevant to the other—captures relationships typically studied in the network games literature. In practice, we should expect the two networks G and $\tilde{G}$ to have many links in common. If G is sufficiently dense, all players know the strategies of those who are payoff-relevant, and local information concerning strategies is sufficient to ensure Nash play. At the opposite extreme, if players possess no personally useful strategic information, then their strategic information does nothing to restrict play.[8]

PROPOSITION 2: *In a simultaneous-move game, the set of peer-confirming equilibria R_G satisfies the following*:

(a) *If $j \notin G_i$ whenever player j is payoff-relevant to player i, then R_G is the set of* (*correlated*) *rationalizable strategy profiles*.

(b) *If $j \in G_i$ whenever player j is payoff-relevant to player i, then R_G is the set of Nash equilibria*.

PROOF: See the Appendix. *Q.E.D.*

Proposition 2 mirrors the first two parts of Proposition 1. Part (a) tells us that, absent any connections to payoff-relevant players, we cannot ensure coordination beyond what rationality and common belief of rationality imply. We may as well have an empty network. Part (b) says that, if only neighbors matter for incentives, peer-confirming equilibrium admits no miscoordination. While a player with few neighbors faces much strategic uncertainty, players being connected to the "right" people is ultimately what matters for attaining coordinated outcomes.

Given the above proposition, a natural conjecture is that the only links in G that matter are those for which at least one player is payoff-relevant to the other, that is, elements of $G \cap \tilde{G}$. This conjecture is false, as we demonstrate at the end of this section through our "Follow the Leader" example.

2.1. *Examples*

We present examples to clarify the definition and illustrate how the set of peer-confirming equilibria varies with G. Even when G is insufficiently dense to obtain Nash equilibrium play, the connections that do exist limit the kinds of miscoordination that can occur. Individual players may benefit from being more connected in G because of the strategic information they possess.

Crowding Out Investment

Three players have an investment opportunity subject to crowding out effects. The players simultaneously decide whether to invest or not. A player who invests (action I) incurs

[8]This may be surprising given that the graph can still be highly connected. The intuition is that player i cannot make useful inferences about j's other neighbors from j's play, unless those neighbors are relevant *to player j*.

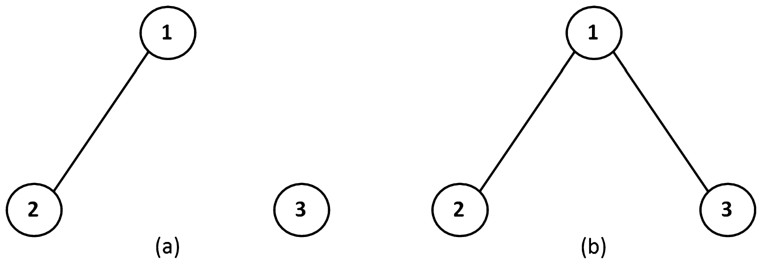

FIGURE 1.—Two possible networks in the investment game.

a positive cost $c \in (\frac{1}{2}, 1)$. If at least one player invests, a unit surplus is generated, and this is divided evenly among those who invest. All players who choose not to invest (action N) earn a default payoff of zero. There are three pure strategy Nash equilibria—(I, N, N), (N, I, N), and (N, N, I)—corresponding to each possible permutation of one player investing while the others refrain. All eight pure strategy profiles are rationalizable: significant miscoordination can occur if players do not accurately anticipate each other's strategies.

Now consider peer-confirming equilibria for the two graphs in Figure 1. In graph (a), players 1 and 2 are linked, while player 3 has no connections. The permissible pure strategy profiles are then (I, N, N), (I, N, I), (N, I, N), (N, I, I), (N, N, I), and (N, N, N). The first two players never invest together, but all other inefficient outcomes are possible. Player 3 may crowd out the investment of one of the first two players, or no one may invest because every player expects someone else to do so.

When we add a link between players 1 and 3 as in graph (b), this further reduces the set of permissible profiles. The peer-confirming equilibria are (I, N, N), (N, I, N), (N, N, I), and (N, I, I). Now, the only possible miscoordination is between players 2 and 3. This network structure confers an "advantage" to player 1, as she is the only player who is guaranteed a nonnegative payoff. Moreover, the additional link results in an important qualitative difference in the outcome set: someone always invests. If neither player 2 nor player 3 plans to invest, player 1 knows this and therefore invests.

Follow the Leader

Three players face a binary choice between action 0 and action 1. Player 1 is indifferent between the two actions, and each of player 2 and player 3 earns a positive payoff if and only if she matches the action of player 1. In a pure strategy Nash equilibrium, all players choose the same action; any profile is rationalizable.

This example helps illustrate how local information can lead to global consequences. Suppose player 2 is linked to both other players, but player 3 is not linked to player 1. In a peer-confirming equilibrium, it is clear that 2 must match 1: player 2 correctly anticipates player 1's strategy, so she chooses the action to match. Player 3 correctly anticipates the action that player 2 is taking, but this action is not relevant to player 3's payoff. Since 3 lacks a link to 1, we might think that 3 can entertain any conjecture about player 1, and therefore does not necessarily match. This naïvete is ruled out by peer-confirming equilibrium. Player 3 is certain that 2 will match 1, and since she correctly anticipates 2's behavior, she does the same for 1's. Nash equilibrium therefore arises, despite the missing link in the network. Under our solution concept, players make sophisticated inferences about those to whom they are not adjacent based on the choices of those to whom they are.

Notice that the above example features a link between two players (2 and 3) who are mutually payoff-irrelevant. If we remove this link, the set of PCE strategy profiles strictly grows to include every profile in which 1 and 2 match. Therefore, in spite of Proposition 2, a peer who exerts no payoff externality can still be an important source of strategic information. Such transmission of strategic information again becomes impossible if no player is *even indirectly* connected to a payoff-relevant player.[9]

3. COORDINATION AND CONNECTIVITY

Players with many connections can often facilitate coordination by others. We say that a player is **fully connected** if she is adjacent to every other player in G. Fully connected players necessarily best respond to the true strategy profile. Therefore, the choice of a fully connected player has the potential to transmit a great deal of strategic information. Whether this actually occurs is sensitive to details of the game. In this section, two examples highlight how a fully connected player's role changes with the structure of a game. If the action space is rich, and best responses are appropriately sensitive to others' actions, a fully connected player can lead to Nash equilibrium play. With coarser actions, when best responses do less to identify others' play, a fully connected player has relatively little impact on the set of peer-confirming equilibria.

3.1. *Public Goods Provision*

Each player i in a population of size N chooses how much to invest in a public good $x_i \in [0, 1]$. The payoff to player i is the benefit from the total investment less the cost of her own investment:

$$u_i(x) = 2\sqrt{\sum_{j=1}^{N} x_j} - x_i.$$

The marginal benefit of i's own investment is

$$\frac{1}{\sqrt{\sum_{j=1}^{N} x_j}} - 1,$$

which immediately implies that in any Nash equilibrium, total investment is exactly 1; conversely, any such (pure) strategy profile is a Nash equilibrium. Player i can rationalize any investment level $x_i \in [0, 1]$ if she expects other players' investments to satisfy $\sum_{j \neq i} x_j = 1 - x_i$. Hence, the range of total investment that can appear in a rationalizable strategy profile is $[0, N]$.

We characterize precisely how the range of total investment in a peer-confirming equilibrium varies with the structure of the graph G. The upper bound depends on the largest *independent set* of players, while the lower bound depends on the presence or absence of a fully connected player.

[9] Formally, suppose $\hat{G} \subseteq G$ is such that every link in $G \setminus \hat{G}$ belongs to a G-connected component whose players are pairwise payoff-irrelevant. Then (as can be shown by adapting the proof of the first part of Proposition 2), we will always have $R_{\hat{G}} = R_G$. We thank an anonymous referee for this observation.

DEFINITION 2: A subset M of players is **independent** if no two players in M are adjacent in G (i.e., $j \notin G_i$ for all $i, j \in M$). An independent subset M is **largest** if there is no independent subset M' with $|M'| > |M|$.

PROPOSITION 3: *Let M denote a largest independent subset of players in G. In the public goods provision game*:

(a) *The maximal investment in any peer-confirming equilibrium is $|M|$.*

(b) *If there exists a fully connected player in G, then the minimal investment in any peer-confirming equilibrium is* 1. *Otherwise, the minimal investment is* 0.

PROOF: First note that a strategy profile $x \in [0, 1]^N$ is a PCE if and only if, for each $i \in N$:

- $x_i \leq 1 - \sum_{j \in G_i} x_j$ if $x_i > 0$;
- $\sum_{j \in N} x_j \geq 1$ if i is fully connected.

The above conditions are necessary for a network-consistent best reply by player i, and therefore for a PCE. To see they are sufficient, consider two cases. If there is a fully connected player who invests, then $\sum_{i \in N} x_i = 1$, so that x is a Nash equilibrium. Suppose no fully connected player invests; every other player i has some non-neighbor $\tilde{i}$. We can support the profile x as a PCE with each non-fully connected $i \in N$ believing the play is a Nash equilibrium with $\tilde{i}$ investing the residual $1 - x_i - \sum_{j \in G_i} x_j$ (and any fully connected players having correct beliefs).

Next, we show that every PCE total investment level is attainable with no two adjacent players investing. Indeed, if i and some neighbor both invest a positive amount, the adjusted profile x' given by

$$x'_j := \begin{cases} x_i + \sum_{k \in G_i} x_k & : j = i, \\ 0 & : j \in G_i, \\ x_j & : \text{otherwise}, \end{cases}$$

attains the same level of total investment in PCE, with strictly fewer investing players. Iterating this construction proves the claim.

Combining the above two claims, X is a PCE total investment level if and only if there is an independent set M of players such that:

- $0 \leq X \leq |M|$;
- if some $i \in N$ is fully connected, then $X \geq 1$.

Both parts of the proposition follow directly. *Q.E.D.*

Proposition 3 gives us insight not only into the range of possible investment levels, but also into the distribution of investment across the players. The presence of a fully connected player ensures that at least the Nash equilibrium level of investment is made, but if such a player actually invests, then no more is invested. If a fully connected player invests a positive amount, it signals the true marginal value of investment to all other players, leading to Nash equilibrium play. When fully connected players do not invest, only then can we obtain a high level of total investment. When this occurs, the fully connected players free-ride while many peripheral players mistakenly believe that investment is their responsibility.

We can also use peer-confirming equilibrium to study other variants of the public goods game. For instance, we might consider the "best-shot" utility

$$u_i(x) = 2\sqrt{\max_{j\in N} x_j} - x_i,$$

or the "weakest-link" utility

$$u_i(x) = 2\sqrt{\min_{j\in N} x_j} - x_i.$$

Some results for the best-shot public goods game are similar to those already presented. Two adjacent players will never both choose positive investment, since at least one of them (one who is investing weakly less) would rather invest nothing. The set of players investing is thus an independent set, so that the same upper bound on total investment holds. This bound is tight, as each player in an independent set could think she is the sole investor in a Nash equilibrium. By an identical argument to that above, the minimum total investment level is zero if there is no fully connected player. Characterizing minimum investment when there is a fully connected player is more delicate in the best-shot public goods model, as non-investment by a fully connected player will provide strategically meaningful information about other players (namely, that at least one player is investing at least $\frac{1}{4}$, so that the fully connected player is willing to not invest).

The weakest-link game produces a qualitatively different set of outcomes, which is perhaps unsurprising since it features strategic complements rather than strategic substitutes. If i and j are adjacent, then $x_i \leq x_j$ in any network-consistent best reply for i. Symmetry and induction imply that players in the same connected component of G must choose the same investment level in any peer-confirming equilibrium. Players could coordinate on any level in $[0, 1]$, but a connected graph implies a Nash equilibrium outcome. If there are multiple connected components, each component may coordinate independently on any investment level in $[0, 1]$—supported by the potentially incorrect belief that the other components choose the same investment level.

3.2. *The Potluck Game*

Each of N players plans to attend a potluck event, and each will bring one of N possible dishes. Players have a common utility function that is monotonically increasing in the number of distinct dishes brought to the event. In any pure strategy Nash equilibrium, all players bring a different dish, resulting in N distinct dishes at the event. Any strategy profile is rationalizable, and players could end up with N pots of the same dish. What networks will facilitate having a higher number of distinct dishes?

We approach this question looking at worst-case outcomes: what is the smallest number of dishes that can appear in a peer-confirming equilibrium? In any profile, two linked players must bring different dishes, and Nash beliefs can support any such profile as a peer-confirming equilibrium. Consequently, the question of how many dishes are brought becomes a question about graph colorings. Recall that the chromatic number of a graph is the minimal number of colors required to color each vertex so that no two adjacent vertices have the same color. The following proposition is immediate from this definition.

PROPOSITION 4: *The minimal number of dishes in a peer-confirming equilibrium is the chromatic number of G.*

Contrasting this with the results for public goods provision reveals how peer-confirming equilibrium is sensitive to the details of a particular game. Fully connected players have a far less significant effect on the outcome set in the potluck game than in the public goods provision game. A fully connected player has complete information about the strategy profile being played. We can guarantee equilibrium play if the fully connected player's best response provides a sufficient statistic for the entire strategy profile. For instance, in the public goods provision game, the marginal value of investment is a sufficient statistic, and the fully connected player's action conveys this marginal value whenever she makes a positive level of investment. In the potluck game, a star network is nearly as bad as the empty network for coordination (with as few as two dishes being served) because the fully connected player's action conveys little useful information to the other players.

4. POLITICAL PROTEST AND ELITE COORDINATION

Successful protests or coups require coordination among many people, and mistakes are costly. Miscoordination can result in prison time or worse for participants in failed efforts, but there may also be high potential gains. When are protests most likely to occur? When are they likely to succeed? When is there a high risk of miscoordination? Which types of miscoordination might we expect? Such questions are often studied using binary action coordination games. Peer-confirming equilibrium can offer a new perspective, allowing us to isolate the role of strategic uncertainty in this setting.

We study a simple binary action coordination game: players move simultaneously, payoffs are symmetric, and there is complete information. Each player in a population of finite size N chooses whether to protest to effect political change. If at least $M \geq 2$ players protest, the regime is overthrown, and each protester earns a payoff $y > 0$. If fewer than M players protest, the protesters are arrested, incurring a cost $c > 0$. Those who do not protest earn a default payoff of 0. Any profile of actions is rationalizable, and there are two pure strategy Nash equilibria: all protest or none protest.

Fully connected players in the graph G are key facilitators of coordination in peer-confirming equilibria. Much like in the public goods provision example, a player who is adjacent to all others conveys valuable information through his choice, and we reach a Nash equilibrium. Even without a fully connected player, however, peer-confirming equilibrium ensures some nontrivial coordination. Suppose there exist two players i and j such that $G_i \cup G_j = N$, that is, players i and j are linked, and every player is linked to at least one of them. We might imagine two distinct communities or political parties with i and j as leaders or representatives. While equilibrium play is not guaranteed, we show that the two leaders must take the same action in any peer-confirming equilibrium.

PROPOSITION 5: *In the political protest game*:

(a) *If there is a fully connected player, then all players choose the same action in any peer-confirming equilibrium*.

(b) *If $G_i \cup G_j = N$, then players i and j choose the same action in any peer-confirming equilibrium*.

PROOF: The first claim is straightforward. The fully connected player protests if and only if the protest will succeed. The other players therefore coordinate on the same action.[10]

[10]Alternatively, the first claim will follow directly from the second claim, taking i to be the fully connected player and letting j range over $N \setminus \{i\}$.

Toward the second claim, suppose $G_i \cup G_j = N$. Without loss, assume that i and j have no neighbors in common, as Proposition 1 tells us this can only expand the set of peer-confirming equilibria. Let Σ denote the set of peer-confirming equilibria. Given a strategy profile σ, write $m_k(\sigma)$ for the number of neighbors of $k \in N$ who protest under σ.

Assume, for a contradiction, that there exists some peer-confirming equilibrium in which exactly one of $\{i, j\}$ protests; without loss, say i protests and j does not. Then define

$$\underline{m}_i := \min\{m_i(\sigma) : \sigma \in R_G, i \text{ protests under } \sigma, j \text{ does not protest under } \sigma\},$$

$$\overline{m}_j := \max\{m_j(\sigma) : \sigma \in R_G, i \text{ protests under } \sigma, j \text{ does not protest under } \sigma\}.$$

As $R_G \subseteq B_G(R_G)$, we can deduce from player i's choice that $\underline{m}_i + \overline{m}_j \geq M - 1$, and from player j's choice that $\overline{m}_j + \underline{m}_i < M - 1$, a contradiction. Hence, the two must choose the same action. *Q.E.D.*

The action of a fully connected player provides a signal allowing others to coordinate. With such a player, we do not necessarily expect more or fewer protests, but the ones that occur are more likely to succeed. After defining peer-confirming equilibrium for dynamic games in the next section, we show that a fully connected player becomes even more significant in facilitating protests if she can publicly commit to an action before other players decide.

If the population is split into two groups, with leaders who interact, we see a particular pattern of miscoordination. The two leaders always move together, but the rest of the players need not. Each leader could protest, with the mistaken belief that enough of the other group will support this action. Alternatively, each leader could refrain, with the mistaken belief that too few in the other group plan to protest. This suggests a mechanism for coordination among elites even as many others choose different actions.

5. PCE IN DYNAMIC GAMES

In this section, we generalize peer-confirming equilibrium to dynamic games. To simplify the analysis, we restrict attention to pure strategies in finite multistage games of observable action.[11] Wherever possible, we mirror the notation already introduced in Section 2.

Consider an extensive-form game $\Gamma = \langle N, H, Z, (A_i^h)_{h\in H, i\in N}, (u_i)_{i\in H}\rangle$. As before, N is a finite set of players. The sets H and Z are the non-terminal and terminal histories, respectively; write $\mathcal{H} = H \cup Z$ for the (finite) set of all histories. For each $h \in H$, the set A_i^h is a finite, nonempty set of actions for player i. We write $A^h = \prod_{i\in N} A_i^h$ for the set of action profiles at history h. Finally, for each $i \in N$, the function $u_i : Z \to \mathbb{R}$ gives player i's utility as a function of the terminal history.

For convenience, we represent elements of $\mathcal{H}$ as finite sequences of elements in $\mathcal{A} = \bigcup_{h\in H} A^h$. We necessarily have $\emptyset \in \mathcal{H}$, and for any $h \in H$ and $a \in \mathcal{A}$, we have $(h, a) \in \mathcal{H}$ if and only if $a \in A^h$. Given two histories $h, h' \in \mathcal{H}$, we write $h \preceq h'$ if h is an initial subsequence of h', that is, $h' = (h, a_1, a_2, \ldots, a_k)$ for some sequence of action profiles $(a_\ell)_{\ell=1}^k$.

[11] Osborne and Rubinstein (1994) refer to such a game as "extensive game[s] with perfect information and simultaneous moves." The definition here is slightly different, though strategically equivalent. We dispense with their player correspondence by allowing players' action sets at a given history to be singletons.

The following notation helps us talk about strategies and their relationship to histories. For each $i \in N$, write $S_i = \prod_{h \in H} A_i^h$ for player i's set of (pure) strategies. A strategy for player i gives the action she takes at any history. For any subset $J \subseteq N$, we similarly write $S_J = \prod_{j \in J} S_j$ for the strategies of players in J, and we define $S \equiv S_N$. For any $h \in \mathcal{H}$ and $J \subseteq N$, we write $S_J(h)$ for the set of partial strategy profiles $s_J \in S_J$ that allow history h to be reached. Each profile in $S_J(h)$ prescribes the same behavior up to history h and arbitrary behavior thereafter and at histories incompatible with h. Similarly, for any $J \subseteq N$ and $s_J \in S_J$, write $H(s_J) = \{h \in H : s_J \in S_J(h)\}$ for the set of non-terminal histories that s_J allows. The set $H(s_J)$ is the set of non-terminal histories that are potentially reached when players in J play according to s_J. Finally, for any $s \in S$ and $h \in H$, we write $\zeta^h(s) \in Z$ for the unique terminal history $\succeq h$ that corresponds to the strategy profile s being followed starting from history h, and $u_i^h(s) := u_i \circ \zeta^h(s)$.

Next, we give notation to formalize players' beliefs and best replies. A **conditional probability system (CPS)** for player i is a vector of distributions $\mu_i = (\mu_i^h)_{h \in H} \in [\Delta(S_{-i})]^H$ that satisfies the following properties:

(a) For all $h \in H$, we have $\mu_i^h(S_{-i}(h)) = 1$.

(b) If $h \preceq h'$ and $E \subseteq S_{-i}(h')$, we have

$$\mu_i^h(E) = \mu_i^h\big(S_{-i}(h')\big)\mu_i^{h'}(E).$$

We interpret μ_i^h as player i's beliefs about opponents' strategies at history h. Condition (a) states that the player assigns probability 1 to what she has already observed: if we are at history h, the other players' strategies must allow history h to be reached. Condition (b) states that player i updates her beliefs according to Bayes's rule as play progresses. The strategy $s_i^* \in S_i$ is a **sequential best reply** to μ_i if, for all $h \in H$, we have

$$s_i^* \in \arg\max_{s_i \in S_i} \int_{S_{-i}} u_i^h(s_i, \cdot)\, \mathrm{d}\mu_i^h.$$

In words, at every history, it maximizes player i's expected payoff going forward. As before, we write $r_i(\mu_i) \subseteq S_i$ for player i's set of sequential best replies to μ_i.

We represent the network of connections between players as an undirected graph G, yielding the augmented game $\langle \Gamma, G \rangle$. For any $j \in N$, $\sigma \in S$, and $h \in H$, define

$$S_j^\sigma(h) = \big\{s_j \in S_j : s_j\big(h'\big) = \sigma_j\big(h'\big)\ \forall h' \in H \text{ with } h' \succeq h\big\}.$$

The set $S_j^\sigma(h)$ is just the set of strategies for player j that agree with σ going forward from h. For each $i \in N$, $\sigma \in S$, and $h \in H$, the set of **continuation plays consistent with i's strategic information at σ** is

$$S_{-i}^{\sigma,G}(h) = \big\{s_{-i} \in S_{-i} : s_j \in S_j^\sigma(h)\ \forall j \in G_i\big\}.$$

Starting from history h, this is the set of opponents' strategies in which all players in G_i follow σ going forward. As we did for simultaneous-move games, we introduce notation for play in $S_{-i}^{\sigma,G}(h)$ that is compatible with $\Sigma \subseteq S$: for each $i \in N$, $\sigma \in S$, $\Sigma \subseteq S$, and $h \in H$, we define

$$\Sigma_{-i}^{\sigma,G}(h) = S_{-i}^{\sigma,G}(h) \cap \big\{s_{-i} \in S_{-i} : (\sigma_i, s_{-i}) \in \Sigma\big\}.$$

We also simply write $\Sigma_{-i}^{\sigma,G}$ in place of $\Sigma_{-i}^{\sigma,G}(\emptyset)$ for the set of opponent strategy profiles consistent with Σ and i's initial strategic information at σ.

Given a strategy profile $\sigma \in S$ and a set of strategy profiles $\Sigma \subseteq S$, for each player $i \in N$ we define the set of **viable conjectures relative to Σ at** σ, denoted $\Delta_i^{\sigma,G}(\Sigma)$. This set consists of all conditional probability systems μ_i such that, for all $h \in H$:

(a) $\mu_i^h(S_{-i}^{\sigma,G}(h)) = 1$;

(b) if $\Sigma_{-i}^{\sigma,G} \cap S_{-i}(h) \neq \emptyset$, then $\mu_i^h(\Sigma_{-i}^{\sigma,G}(h)) = 1$.

Condition (a) requires that, at each history $h \in H$, player i's beliefs assign probability 1 to the event that i's neighbors follow σ going forward. Condition (b) states that, if some strategy consistent with Σ and i's strategic information at σ allows us to reach history h, then i's beliefs at history h assign probability 1 to $\Sigma_{-i}^{\sigma,G}$. This means that, if player i starts the game believing the strategy profile is in $\Sigma_{-i}^{\sigma,G}$, she must maintain this belief throughout the game as long as she observes no conclusive proof to the contrary. This is sometimes referred to as strong belief.

Using the set of viable conjectures, we can define the **network-consistent best replies** to Σ:

$$B_G(\Sigma) = \{\sigma \in \Sigma : \forall i \in N, \exists \mu_i \in \Delta_i^{\sigma,G}(\Sigma) \text{ s.t. } \sigma_i \in r_i(\mu_i)\}.$$

The set $B_G(\Sigma)$ includes all strategy profiles in Σ for which every player follows a sequential best reply to some viable conjecture. We can now state our definition of peer-confirming equilibrium in dynamic games.

DEFINITION 3: The set of **peer-confirming equilibria** for $\langle \Gamma, G \rangle$ is

$$R_G \equiv \bigcap_{k=0}^{\infty} B_G^k(S).$$

A profile $\sigma \in S$ is a peer-confirming equilibrium if $\sigma \in R_G$.

Analogously to Proposition 1, the set of peer-confirming equilibria for dynamic games has the following properties.

PROPOSITION 6: *In a dynamic game, the set of peer-confirming equilibria R_G satisfies the following*:

(a) *If G is empty, then R_G is the set of extensive-form rationalizable (a.k.a. strongly rationalizable) strategy profiles.*[12]

(b) *If G is complete, then R_G is the set of (pure strategy) subgame perfect equilibria.*

(c) *The set of peer-confirming equilibria of Γ is contained in the set of peer-confirming equilibria for the normal form of Γ (the latter being viewed as a simultaneous-move game).*

PROOF: See the Appendix. *Q.E.D.*

While simple, Proposition 6 already reveals nuance to peer-confirming equilibrium in dynamic games. Peer-confirming equilibrium specializes to extensive-form rationalizability in the empty network and to subgame perfect equilibrium in the complete network. In general, these are non-nested solution concepts. Hence, we can obtain no analogue of Proposition 1(c) for dynamic games: the peer-confirming equilibrium set is

[12]The literature exhibits several slightly different definitions of extensive-form rationalizability. See the discussion preceding the proposition's proof in the Appendix for clarification.

non-monotonic in the network. This is due to the notion of strong belief—narrowing the set of beliefs a player might entertain can limit the refining power of forward induction (Battigalli and Friedenberg (2012)). Far from a technical complication, this feature adds economic content to peer-confirming equilibrium. As we see in a dynamic protest game in the following section, forward induction reasoning and local strategic information can interact, refining away some equilibrium outcomes. In spite of this, the following result shows that equilibria which are unambiguously good for all decision makers will never get refined away.

PROPOSITION 7: *If $\sigma \in S$ has the property that $\sigma \in \arg\max_{s \in S(h)} u_i(s)$ for every $h \in H$ and $i \in N$ such that $|A_i^h| > 1$, then σ is a peer-confirming equilibrium for every network G.*

PROOF: See the Appendix. *Q.E.D.*

6. SIGNALING OTHERS' INTENTIONS

Applying our solution concept yields new insights in a dynamic version of the political protest game (see Section 4 for a static analogue). As before, there is a population of N players, each of whom makes a choice whether to protest. If at least $M \in [3, N-1]$ players protest, each protester earns $y > 0$. If fewer than M protest, each protester then incurs the cost $c > 0$. Unlike the earlier model, players do not move simultaneously. One player is a "leader" and makes a publicly observable choice before anyone else. After the leader decides whether to protest, the other players make simultaneous choices.

We contrast the sets of peer-confirming equilibria for two networks: a complete network and a star network centered on the leader (see Figure 2). Proposition 6 implies that in the former case, we obtain the set of pure strategy subgame perfect equilibria. In particular, the complete network admits equilibria in which a protest is successful and equilibria in which there is no protest.

In the star network, we actually get a *smaller* set of outcomes. To understand why, think about the reasoning of the other players after observing the leader commit to protest. The leader is adjacent to all others in the star network, implying she has an accurate conjecture about everyone's behavior. If the leader is rational, she only commits to protest if she expects at least $M - 1$ others to follow. Therefore, the other players who observe the leader protest believe that the protest will succeed, so the unique best response for all players is to protest. Knowing this, the leader can *always* induce others to protest by committing herself to do so. Hence, the only peer-confirming equilibrium in the star network is for all players to protest.

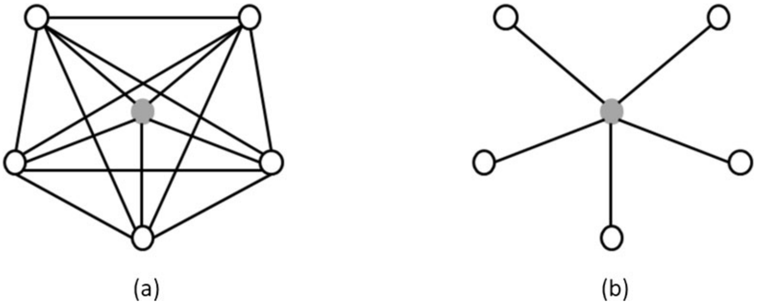

FIGURE 2.—A complete network (a) and a star network (b), leader in gray.

PROPOSITION 8: *In the dynamic protest game with a leader-centered star network, all players protest in peer-confirming equilibrium.*

PROOF: First, by Proposition 7, the strategy profile in which everybody protests at every history is a peer-confirming equilibrium. Therefore, $\Sigma^k := B_G^k(S)$ is nonempty for every $k \in \mathbb{Z}_+$.

Let $\ell \in N$ denote the leader. For all $J \subseteq N \setminus \{\ell\}$ and $m \in \mathbb{Z}_+$, let $P_J(m) \subseteq S_J$ be the set of partial strategy profiles with the feature that, if the leader protests, then at least m people from player subset J will protest as well.

The leader's best response property (together with the fact that the leader is adjacent to all other players) tells us that $\Sigma^1 \cap (\{\text{NoProtest}\} \times P_{N\setminus\{\ell\}}(M-1)) = \emptyset$. That is, if an attempted protest by the leader would succeed, then the leader will indeed protest.

Therefore, $\forall i \in N \setminus \{\ell\}$, $\sigma \in \Sigma^1$, $\mu_i \in \Delta_i^{\sigma,G}(\Sigma^1)$:

$$\begin{aligned} &\mu_i^{\text{Protest}}\big(\{\text{Protest}\} \times P_{N\setminus\{\ell,i\}}(M-2)\big) = 1 \\ \Longrightarrow\ &\sigma_i(\text{Protest}) = \text{Protest}, \quad \text{if } \sigma_i \in r_i(\mu_i). \end{aligned}$$

That is, $i \neq \ell$ infers from a leader protest that a protest would succeed (at least with i's support), and therefore optimally protests herself.

In summary, $\forall \sigma \in \Sigma^2$, $\forall i \in N \setminus \{\ell\}$, $\sigma_i(\text{Protest}) = \text{Protest}$. The leader's best response property then tells us that the leader protests in any $\sigma \in \Sigma^3$. As $\Sigma^3 \subseteq \Sigma^2$, it follows that everybody protests on path in any $\sigma \in \Sigma^3$, and therefore in any peer-confirming equilibrium. *Q.E.D.*

Why does the above reasoning fail in the complete network? If the leader protests when no one else plans to do so, then the other players—who correctly anticipate each other's strategies—can deduce from their strategic information that the leader erred. In the star network, the other players have no strategic information to contradict their belief in the leader's rationality, so they must conclude that a protest will succeed. This highlights the subtle role that strategic information can play in facilitating coordination. In some cases, having less information can help players coordinate on a particular outcome because the past choices of some players become more meaningful signals of the future behavior of others. In the context of our application, peer-confirming equilibrium illustrates the informative role of leadership in revolutions (Shadmehr and Bernhardt (2017)).

This example also highlights that extensive-form PCE delivers a novel form of forward induction. Under forward induction reasoning, "surprise events are regarded as arising out of purposeful choices of the opponents," so that "a player may try to draw inferences about future play from a past surprising choice made by an opponent."[13] In all previous applications of which we are aware, the related inferences always concern the future behavior of the player whose choice was surprising. That is, past behavior of player i has no predictive power for the future behavior of player $j \neq i$; in particular, forward induction reasoning is irrelevant if no one moves twice along any path. Under peer-confirming equilibrium, observed behavior of player i can help predict—using the *joint* identifying assumption that i is rational and has correct beliefs about her neighbors' play—future choices of player i's neighbors. As our protest example shows, this ability to signal strategic information can affect real outcomes of a game, even when players move once.

[13] Dekel and Siniscalchi (2015, p. 676).

Chwe (2000) offers an alternative perspective on the role of social networks in coordinating protests. In his paper, players are uncertain whether others are "willing" to protest (i.e., have payoff type which leaves protest undominated), and links in a communication network allow players to inform others about their types. The focus is on "sufficient networks" in which it is possible for protest to happen in equilibrium—even when all players are willing, uncertainty about others' types may render protest impossible. In that paper, having more connections always facilitates protest, to which uncertainty about others' payoff types is the main impediment. Minimal sufficient networks have a nested structure in which a clique of leaders is able to share information with everyone else. Our protest example features no payoff uncertainty: all players are "willing" to protest with probability 1. Rather than focus on whether protest is *possible*, our example explores whether it is *necessary*. Here, there is no direct communication, but the leader's action conveys information beyond her own plans.

7. RELATIONSHIP TO OTHER SOLUTION CONCEPTS

This section explores the relationship between peer-confirming equilibrium and other solution concepts. In simultaneous-move games, peer-confirming equilibrium becomes a special case of rationalizable conjectural equilibrium (RCE) in which the network determines the corresponding information partition. In simultaneous-move games, peer-confirming equilibrium is likewise a special case of rationalizable partition-confirmed equilibrium (RPCE). However, in dynamic games, our forward induction refinement places restrictions on off-path beliefs that are not present in RPCE, leading to fundamental differences.

7.1. *Rationalizable Conjectural Equilibrium*

Rationalizable conjectural equilibrium is defined in static settings (Rubinstein and Wolinsky (1994)).[14] Given a realized strategy profile s, each player i observes an associated *signal* $g_i(s)$. Each player's signal induces a partition of the space of strategy profiles; informally, players can only distinguish between outcomes if they are in different partition elements.

Just as the network is transparent to the players in (the interpretation implicit in) our definition, the signal maps are transparent to the players under RCE. In RCE, players are rational, players correctly anticipate their own realized signals, and these facts are commonly believed. Player i forms a conjecture over the choices s_{-i} of other players and plays a best response. In equilibrium, this conjecture must put probability 1 on the same realized signal as the true action profile. If g_i is one-to-one for each player i, the concept reduces to (pure strategy) Nash equilibrium. If g_i is constant for each player i, the concept reduces to (correlated) rationalizability.

Peer-confirming equilibrium is a special case that imposes additional structure—the network determines players' signals. Setting $g_i(s) = (s_j)_{j\in G_i}$, the definition of RCE exactly reduces to a PCE. Since the actions of neighbors are fully identified by the signal, conjectures must assign probability 1 to neighbors' true actions. Note that, like in RCE, our solution concept allows correlated conjectures: a conjecture is a distribution on the set of opponent action profiles.

[14] RCE is defined on the normal form of a game. For ease of comparison, we interpret it here in the context of simultaneous-move games.

7.2. *Rationalizable Partition-Confirmed Equilibrium*

In RPCE, each player has correct beliefs over a partition of the possible terminal nodes in the game. As Fudenberg and Kamada (2015) observe, in a simultaneous-move game, rationalizable conjectural equilibrium and rationalizable partition-confirmed equilibrium are equivalent, because the set of terminal nodes is simply the set of action profiles. Consequently, in a simultaneous-move game, peer-confirming equilibrium specializes RPCE.

In dynamic games, this relationship no longer holds. The key difference lies in restrictions on off-path beliefs. Fudenberg and Kamada (2015) state their restrictions in their definition of *accordance*: on the equilibrium path, players update according to Bayes's rule; off path, a player assigns probability 1 to profiles supported in her prior. Given the mildness of the belief restriction, *every* subgame perfect equilibrium is an RPCE, whatever the players' partitions. In a peer-confirming equilibrium, players are forced to rationalize their off-path beliefs.

Consider the two-stage protest game of the previous section. Mapping the star network to a terminal node partition, the leader gets to observe the terminal node exactly, while the other players observe only the leader's action in addition to their own. We can trivially obtain no protest in a RPCE with the leader correctly anticipating that no one will follow a choice to protest, and all other players believing that no one else will protest, regardless of which of the two possible information sets they observe. If the leader deviates, nothing forces the other players to abandon their belief that the protest will fail.

Theorem 2 in Fudenberg and Kamada (2015) further highlights the difference with peer-confirming equilibrium. The set of RPCE is monotone with respect to the fineness of the terminal node partition. When the partition is determined by a network, this means that the set of RPCE is monotone in the network: adding links weakly shrinks the set of RPCE. As we saw in the previous two sections, PCE does not enjoy the same monotonicity property.

8. FINAL REMARKS

Peer-confirming equilibrium gives us an intuitive and practical way to use the structure of social relationships to refine outcome predictions in games. We typically have clearer expectations about the behavior of our friends, neighbors, and closest colleagues than about others, and peer-confirming equilibrium acknowledges that this should influence model predictions. Through this concept, we gain insights about who in a network is important for coordination—fully connected individuals are often pivotal, but this finding is sensitive to details of the underlying game—and how a game's payoff structure interacts with strategic information. These insights suggest testable propositions for future study in empirical work.

We also highlight a new channel through which actions convey information. When knowledge about players' strategies is incomplete, one player's action can provide a signal about others' choices. In a static game, the signal is implicit: knowing a player's intended action entails knowing something about her beliefs about others' strategies. Under the right circumstances, the behavior of a fully connected player can effectively serve as a sufficient statistic for everyone's behavior, enabling coordination on equilibrium. As we saw in the protest example, even without a fully connected player, we can draw nontrivial conclusions about who coordinates with whom. In dynamic games, signaling is more explicit, though its effects are more intricate. Our solution concept sometimes leads to sharper predictions than either extensive-form rationalizability or subgame perfect equilibrium, which may make it well-suited for applied work.

The intuition behind peer-confirming equilibrium clearly suggests broader application to games in which a network also captures payoff relationships. In some of these cases, the usefulness of our concept is unclear. For instance, network formation games typically use cooperative solution concepts like pairwise stability because noncooperative concepts produce an enormous multiplicity (Jackson and Wolinsky (1996)). In other games, particularly those involving local strategic complements and substitutes (e.g., Ballester, Calvó-Armengol, and Zenou (2006), Bramoullé and Kranton (2007)), applications look more promising.

We see potential applications of our concept to the study of organizations and communication games. Empirically, coordination on a common understanding of linguistic conventions is a nontrivial problem (e.g., Weber and Camerer (2003), Selten and Warglien (2007)), and PCE offers one way to formally describe who understands whose messages. Hence, we believe PCE may serve as a useful modeling ingredient to study the value of investments in codes and other communication resources. Another possible application is to use PCE to endogenize the network itself. In conjunction with a model of network formation, it is tempting to consider a richer extension of peer-confirming equilibrium in dynamic games with different networks obtaining at different points in the game. Whether such an extension adds any meaningful richness to the strategic analysis remains to be seen.

In many network games, measures of player centrality—degree centrality, Katz–Bonacich centrality, eigenvector centrality—characterize equilibrium behavior. Such centrality measures are conspicuously absent from the characterizations in our examples. We rely instead on an alternative set of graph concepts—connectedness, independent sets, graph colorings—that rarely appear in economic network models. This highlights a fundamental difference in the way that payoffs and strategic information propagate in networks. This may be partly driven by the all-or-nothing nature of links in our framework—one either knows another player's strategy or not. With an appropriate notion of tie strength that captures how well one knows another player's strategy, more traditional centrality measures may emerge, and we feel this presents an intriguing agenda for future work.

A strong assumption which is implicit in the definition of peer-confirming equilibrium is that players know the network structure. For some applications, non-local network knowledge seems realistic, for example, voters may anticipate that politicians of the same state from the same party will face little strategic uncertainty concerning each other. However, the assumption seems unreasonable for other applications. Generalizing peer-confirming equilibrium to incorporate beliefs about the underlying network seems like a natural next step. One concern for extending the model in this way is that beliefs about network structure can be complex. However, our examples offer hope that, at least for some settings, the PCE set depends only on simple network features (e.g., chromatic number, existence of a fully connected player). This may lend tractability to models of incomplete network information, requiring the players (and analyst) only to track such low-dimensional network features.

APPENDIX A: OMITTED PROOFS

Proof of Proposition 2

PROOF: Toward the first part, we first apply Proposition 1 to see that $R_G \subseteq \Sigma$, where $\Sigma \subseteq S$ is the set of correlated rationalizable profiles. It therefore suffices to show that $\Sigma \subseteq B_G(\Sigma)$. To that end, consider any $\sigma \in \Sigma$ and $i \in N$. By definition of rationalizability,

Σ is a product set, and there exists some $\mu_i \in \Delta(\Sigma_{-i})$ such that $\sigma_i \in r_i(\mu_i)$. Now, define the modified belief,

$$\mu_i^\sigma := (\mathrm{marg}_{N\setminus[\{i\}\cup G_i]}\,\mu_i) \otimes \bigotimes_{j\in G_i} \delta_{\sigma_j}.$$

By construction, $\mu_i^\sigma \in \Delta_i^{\sigma,G}(\Sigma)$. By hypothesis, no member of G_i is payoff-relevant to i, so that $\int_{S_{-i}} u_i(s_i,\cdot)\,\mathrm{d}\mu_i^\sigma = \int_{S_{-i}} u_i(s_i,\cdot)\,\mathrm{d}\mu_i$ for every $s_i \in S_i$. In particular $r_i(\mu_i^\sigma) = r_i(\mu_i) \ni \sigma_i$. As $\sigma \in \Sigma$ and $i \in N$ were arbitrary, $\Sigma \subseteq B_G(\Sigma)$, as required.

Toward the second part, suppose σ is a peer-confirming equilibrium. From the definition, for each player i, there is a conjecture μ_i, assigning probability 1 to σ_j for each $j \in G_i$, such that

$$\sigma_i \in \arg\max_{s\in S_i} \int_{S_{-i}} u_i(s, s_{-i})\,\mathrm{d}\mu_i(s_{-i}).$$

From the definition of payoff-relevance, we have

$$\int_{S_{-i}} u_i(s, s_{-i})\,\mathrm{d}\mu_i(s_{-i}) = u_i(s, \sigma_{-i}),$$

since $u_i(s, s_{-i}) = u(s, \sigma_{-i})$ whenever $s_j = \sigma_j$ for each $j \in G_i$. Therefore,

$$\sigma_i \in \arg\max_{s\in S_i} u_i(s, \sigma_{-i}).$$

Consequently, in any profile contained in R_G, players employ best replies to the true strategy profile, implying a Nash equilibrium. Proposition 1 then implies that the peer-confirming equilibrium set and the Nash equilibrium set coincide exactly. *Q.E.D.*

Proof of Proposition 6

For part (a) of the proposition, a remark on terminology is in order. The literature does not exhibit a unified definition of extensive-form rationalizability. Here, we follow Pearce (1984) in imposing rationality of player i's continuation play even at histories precluded by i's own strategy, and we follow Battigalli and Siniscalchi (2002) in allowing for correlated conjectures. Thus, our notion of EFR refines that of Battigalli and Siniscalchi (2002), but in an inessential way: it results in exactly the same admissible *plans of action*, in the sense of Rubinstein (1991).

PROOF: (a) From the definitions, $\forall i \in N$, $\sigma \in S$, and product sets $\Sigma = \prod_{j\in N} \Sigma_j \subseteq S$:

$$\begin{aligned}
& S_{-i}^{\sigma,\emptyset}(h) = S_{-i}, \quad \forall h \in H \\
\Longrightarrow\ & \Sigma_{-i}^{\sigma,\emptyset}(h) = \mathrm{proj}_{-i}\Sigma = \Sigma_{-i}, \quad \forall h \in H \\
\Longrightarrow\ & \Delta_i^{\sigma,\emptyset}(\Sigma) \text{ is the set of CPSs } \mu_i \text{ such that for all } h \in H, \\
& \text{if } \Sigma_{-i} \cap S_{-i}(h) \neq \emptyset, \text{ then } \mu_i^h(\Sigma_{-i}) = 1.
\end{aligned}$$

Therefore, in an empty network, $B_\emptyset$ takes product sets to product sets, and our algorithm specializes to essentially the same algorithm that defines extensive-form rationalizability in Battigalli and Siniscalchi (2002).[15]

(b) Denote the complete network $K = \{ij : i, j \in N, i \neq j\}$.

From the definitions, $\forall i \in N$, $\sigma \in S$, and $\Sigma \subseteq S$:

$$S_{-i}^{\sigma,K}(h) = \prod_{j\neq i} S_j^{\sigma}(h) =: S_{-i}^{\sigma}(h), \quad \forall h \in H$$

$$\implies \Sigma_{-i}^{\sigma,K}(h) = S_{-i}^{\sigma}(h) \cap \{s_{-i} \in S_{-i} : (\sigma_i, s_{-i}) \in \Sigma\}, \quad \forall h \in H$$

$\implies \Delta_i^{\sigma,K}(\Sigma)$ is the set of CPSs μ_i such that for all $h \in H$:

- $\mu_i^h(S_{-i}^{\sigma}(h)) = 1$;
- if $\Sigma \cap [\{\sigma_i\} \times S_{-i}(h)] \neq \emptyset$, then $\mu_i^h\{s_{-i} \in S_{-i} : (\sigma_i, s_{-i}) \in \Sigma\} = 1$.

Applying the above, $\Delta_i^{\sigma,K}(S)$ is the set of CPSs μ_i such that $\forall h \in H$, $\mu_i^h(S_{-i}^{\sigma}(h)) = 1$. Therefore, $\hat{\Sigma} := B_K(S)$ is the set of subgame perfect equilibria of Γ.

Suppose $\sigma \in \hat{\Sigma}$, and consider any $i \in N$. Note that $\Delta_i^{\sigma,K}(\hat{\Sigma}) \neq \emptyset$ as $\hat{\Sigma} \neq \emptyset$ (as witnessed by σ). Then an arbitrary $\mu_i \in \Delta_i^{\sigma,K}(\hat{\Sigma})$ will have $r_i(\mu_i) \ni \sigma_i$ by condition (a) in $\Delta_i^{\sigma,\cdot}$'s definition and the fact that σ is a subgame perfect equilibrium. Therefore, $B_K(\hat{\Sigma}) = \hat{\Sigma}$. It follows that $\hat{\Sigma} = R_K$.

(c) To avoid overloaded notation, we temporarily add the decoration $\tilde{\cdot}$ above any definitions applied to the normal form.

For any $i \in N$ and $\sigma \in \Sigma \subseteq S$:

- $\tilde{S}_{-i}^{\sigma,G} = S_{-i}^{\sigma,G}(\phi) \implies \tilde{\Sigma}_{-i}^{\sigma,G} \supseteq \Sigma_{-i}^{\sigma,G} \implies \tilde{\Delta}_i^{\sigma,G}(\Sigma) \supseteq \Delta_i^{\sigma,G}(\Sigma)$.
- $\tilde{r}_i(\mu_i^{\phi}) \supseteq r_i(\mu_i)$ for every CPS μ_i, and every $\mu_i^{\phi} \in \Delta(S_{-i})$ extends to some CPS μ_i.

Taken together, these tell us that $\tilde{B}_G(\Sigma) \supseteq B_G(\Sigma)$ for any $\Sigma \subseteq S$. Moreover, recall that $\tilde{B}_G$ is monotonic.

Now, if $k \in \mathbb{Z}_+$ and $\tilde{B}_G^k(S) \supseteq B_G^k(S)$, then monotonicity of $\tilde{B}_G$ tells us:

$$\tilde{B}_G^{k+1}(S) = \tilde{B}_G(\tilde{B}_G^k(S)) \supseteq \tilde{B}_G(B_G^k(S)) \supseteq B_G(B_G^k(S)) = B_G^{k+1}(S).$$

By induction, $\tilde{B}_G^k(S) \supseteq B_G^k(S)$, $\forall k \in \mathbb{Z}_+$. Taking intersections, $\tilde{R}_G \supseteq R_G$. *Q.E.D.*

Proof of Proposition 7

PROOF: We first invest in some notation. Given $i \in N$, $s, \sigma \in S$, and $\bar{h} \in H$, define $s_i\bar{h}\sigma_i \in S_i$ via

$$\forall h \in H, \quad s_i\bar{h}\sigma_i(h) := \begin{cases} s_i(h) & : h \not\succeq \bar{h}, \\ \sigma_i(h) & : h \succeq \bar{h}. \end{cases}$$

Similarly, define $s_{-i}\bar{h}\sigma_{-i} := (s_j\bar{h}\sigma_j)_{j\neq i}$ and $s\bar{h}\sigma := (s_j\bar{h}\sigma_j)_{j\in N}$.

[15] "Essentially" because of our slightly refined notion of sequential rationality. As opponent's plans of action remaining at each stage are sufficient to compute own best responses, it is immediate that the set of remaining plans of action is identical to that of Battigalli and Siniscalchi (2002).

Given $\sigma \in S$, say $\Sigma \subseteq S$ is σ-comprehensive if $\forall \hat{\sigma} \in \Sigma$, $\bar{h} \in H(\hat{\sigma})$, we have $\hat{\sigma}\bar{h}\sigma \in \Sigma$, too. That is, Σ is σ-comprehensive if, when we edit a strategy profile from Σ by replacing some on-path continuation play with σ, the edited profile still belongs to Σ.

Claim: If $\Sigma \subseteq S$ is σ-comprehensive, then so is $B_G(\Sigma)$.

Verification: Take $\bar{\sigma} \in B_G(\Sigma)$, $\bar{h} \in H$. We want to show $\sigma^* := \bar{\sigma}\bar{h}\sigma \in B_G(\Sigma)$.

Take any $i \in N$. By hypothesis, $\exists \bar{\mu}_i \in \Delta_i^{\bar{\sigma},G}(\Sigma)$ such that $\bar{\sigma}_i \in r_i(\bar{\mu}_i)$. Define

$$\begin{aligned} \phi : S_{-i} &\to S_{-i}, \\ s_{-i} &\mapsto s_{-i}\bar{h}\sigma_{-i}, \end{aligned}$$

and define $\mu_i^* \in [\Delta(S_{-i})]^H$ via $\mu_i^{*h} := \bar{\mu}_i^h \circ \phi^{-1}$ for all $h \in H$.

Now, $\bar{\mu}_i$ is a CPS such that $\bar{\mu}_i^h(S_{-i}^{\bar{\sigma},G}(h)) = 1$. Therefore, μ_i^* is a CPS such that $\mu_i^{*h}(S_{-i}^{\sigma^*,G}(h)) = 1$ by construction. Next, if $\bar{\mu}_i^h(\Sigma_{-i}^{\bar{\sigma},G}) = 1$ for some $h \in H$, then $\mu_i^{*h}(\Sigma_{-i}^{\sigma^*,G}) = 1$ because Σ is σ-comprehensive. Therefore, $\mu_i^* \in \Delta_i^{\sigma^*,G}(\Sigma)$. Finally, that $\bar{\sigma}_i \in r_i(\bar{\mu}_i)$, together with the sequential first-best hypothesis on σ, implies $\sigma_i^* \in r_i(\mu_i^*)$. So $\sigma^* \in B_G(\Sigma)$, proving the claim.

Claim: If $\Sigma \subseteq S$ is σ-comprehensive and contains σ, then $B_G(\Sigma) \ni \sigma$.

Verification: Take any $i \in N$. As $\Sigma \neq \emptyset$, $\exists \tilde{\mu}_i \in \Delta_i^{\sigma,G}(\Sigma)$.

We now define a map $\psi : S_{-i} \times H \to S_{-i}$.

Given $\tilde{s}_{-i} \in S_{-i}$, define $(s_{-i}^h)_{h \in H} := \psi(\tilde{s}_{-i}, \cdot)$ recursively as follows. $\forall h \in H$:

- If $\exists h_0 \in H$ such that[16] $h_0 \preceq h$ and $h \in H(s_{-i}^{h_0})$, then let $s_{-i}^h := s_{-i}^{h_0}$.
- Otherwise, let $s_{-i}^h := \tilde{s}_{-i}h\sigma$.

Now, define $\mu_i \in [\Delta(S_{-i})]^H$ via $\mu_i^h := \tilde{\mu}_i^h \circ \psi(\cdot, h)^{-1}$.

As $\tilde{\mu}_i \in \Delta_i^{\sigma,G}(\Sigma)$ and Σ is σ-comprehensive, it follows that $\mu_i \in \Delta_i^{\sigma,G}(\Sigma)$, too. By construction, $\mu_i^h(S_i^\sigma(h)) = 1\ \forall h \in H$. Finally, as σ is a subgame perfect equilibrium, it follows that $\sigma_i \in r_i(\mu_i)$. Therefore, $\sigma \in B_G(\Sigma)$, proving the claim.

By induction, the two claims tell us that $\forall k \in \mathbb{Z}_+$, $B_G^k(S)$ is σ-comprehensive and contains σ. Taking intersections, $R_G \ni \sigma$. *Q.E.D.*

APPENDIX B: AN ALTERNATIVE CHARACTERIZATION OF PCE

Here, we provide an alternative characterization of peer-confirming equilibrium in multistage games.

PROPOSITION 9: *For each $k \in \mathbb{Z}_+$, we have $S_G^k = B_G^k(S)$, where $S_G^0 := S$, and, for each $k \in \mathbb{N}$,*

$$S_G^k := \left\{ \sigma \in S : \forall i \in N, \exists \mu_i \in \bigcap_{l=0}^{k-1} \Delta_i^{\sigma,G}\big(S_G^l\big) \text{ s.t. } \sigma_i \in r_i(\mu_i) \right\}.$$

In particular, $R_G = \bigcap_{k=0}^{\infty} S_G^k$.

PROOF: We induct on k; the base case $k = 0$ holds by definition. Suppose $k \in \mathbb{N}$ with $S_G^{k-1} = B_G^{k-1}(S) =: \Sigma$.

[16] Note that, if there are multiple such h_0, then they have the same $s_{-i}^{h_0}$. The construction is thus well-defined.

Take any $\sigma \in S_G^k$. We clearly have $\sigma \in \Sigma$, since $S_G^k \subseteq S_G^{k-1}$. To show $\sigma \in B_G(\Sigma)$, consider any $i \in N$. Since $\sigma \in S_G^k$, there exists $\mu_i \in \bigcap_{l=0}^{k-1} \Delta_i^{\sigma,G}(S_G^l) \subseteq \Delta_i^{\sigma,G}(\Sigma)$ such that $\sigma_i \in r_i(\mu_i)$. Hence, by definition $\sigma \in B_G(\Sigma)$.

Now take any $\sigma \in B_G(\Sigma)$. Since $\sigma \in \Sigma$, given any $i \in N$, there exists $\overline{\mu}_i \in \bigcap_{l=0}^{k-2} \Delta_i^{\sigma,G}(S_G^l)$ such that $\sigma_i \in r_i(\overline{\mu}_i)$, and there exists $\hat{\mu}_i \in \Delta_i^{\sigma,G}(\Sigma)$ such that $\sigma_i \in r_i(\hat{\mu}_i)$. Define the conditional probability system μ_i by

$$\mu_i^h = \begin{cases} \hat{\mu}_i^h & \text{if } \Sigma_{-i}^{\sigma,G} \cap S_{-i}(h) \neq \emptyset, \\ \overline{\mu}_i & \text{if } \Sigma_{-i}^{\sigma,G} \cap S_{-i}(h) = \emptyset. \end{cases}$$

By construction, both $\sigma_i \in r_i(\mu_i)$ and $\mu_i \in \bigcap_{l=0}^{k-1} \Delta_i^{\sigma,G}(S_G^l)$ hold.[17] Hence, $\sigma \in S_G^k$. *Q.E.D.*

The above result, whose proof is a straightforward adaptation of that of Battigalli (1997, Theorem 1), is perhaps helpful in interpreting the assumptions implicit in PCE for dynamic games, and is at least suggestive of an epistemic characterization. The set S_G^1 is the set of strategies compatible with each player responding rationally to her beliefs, and each player being correct and certain (at all histories) concerning the continuation play of her neighbors. For a given $k \in \mathbb{N}$, then, S_G^{k+1} is meant to capture the play compatible with S_G^1, strong belief in $S_G^1, \dots,$ and strong belief in S_G^k. So, informally, we interpret PCE as capturing sequential rationality, full correct conjectures concerning neighbors' future play, and common strong belief thereof.

REFERENCES

ANGELETOS, G.-M., C. HELLWIG, AND A. PAVAN (2007): "Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks," *Econometrica*, 75 (3), 711–756. [568,570]

AUMANN, R., AND A. BRANDENBURGER (1995): "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63 (5), 1161–1180. [570]

BACH, C., AND E. TSAKAS (2014): "Pairwise Epistemic Conditions for Nash Equilibrium," *Games and Economic Behavior*, 85, 48–59. [570]

BALLESTER, C., A. CALVÓ-ARMENGOL, AND Y. ZENOU (2006): "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 74 (5), 1403–1417. [569,586]

BANDIERA, O., AND I. RASUL (2006): "Social Networks and Technology Adoption in Northern Mozambique," *The Economic Journal*, 116, 869–902. [569]

BANERJEE, A., A. G. CHANDRASEKHAR, E. DUFLO, AND M. O. JACKSON (2013): "The Diffusion of Microfinance," *Science*, 341 (6144), 1236498. [569]

BATTIGALLI, P. (1997): "On Rationalizability in Extensive Games," *Journal of Economic Theory*, 74 (1), 40–61. [590]

BATTIGALLI, P., AND A. FRIEDENBERG (2012): "Forward Induction Reasoning Revisited," *Theoretical Economics*, 7 (1), 57–98. [582]

BATTIGALLI, P., AND M. SINISCALCHI (2002): "Strong Belief and Forward-Induction Reasoning," *Journal of Economic Theory*, 106, 356–391. [570,587,588]

——— (2003): "Rationalization and Incomplete Information," *B.E. Journal of Theoretical Economics*, 3 (1), 1–46. [570]

BERNHEIM, D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52 (4), 1007–1028. [570,572]

BRAMOULLÉ, Y., AND R. KRANTON (2007): "Public Goods in Networks," *Journal of Economic Theory*, 135, 478–494. [569,586]

BRAMOULLÉ, Y., A. GALEOTTI, AND B. ROGERS (Eds.) (2016): *The Oxford Handbook of the Economics of Networks*. Oxford University Press. [567]

[17]Notice that μ_i respects the Bayesian property of CPSs because (i) the set of $h \in H$ with $\Sigma_{-i}^{\sigma,G} \cap S_{-i}(h) = \emptyset$ is closed under succession, and (ii) any first transition to such a history is necessarily zero probability under $\hat{\mu}_i$.

CARLSSON, H., AND E. VAN DAMME (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61 (5), 989–1018. [570]
CARRELL, S., B. SACERDOTE, AND J. WEST (2013): "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation," *Econometrica*, 81 (3), 855–882. [569]
CATONINI, E. (2017): "Self-Enforcing Agreements and Forward Induction Reasoning," Technical Report, Higher School of Economics. [570]
CHWE, M. (2000): "Communication and Coordination in Social Networks," *Review of Economic Studies*, 67, 1–16. [570,583,584]
DEKEL, E., AND M. SINISCALCHI (2015): "Epistemic Game Theory," in *Handbook of Game Theory*, Vol. 4, ed. by P. Young and S. Zamir. North-Holland. [570,583]
EDMOND, C. (2013): "Information Manipulation, Coordination, and Regime Change," *Review of Economic Studies*, 80 (4), 1422–1458. [570]
ESPONDA, I. (2013): "Rationalizable Conjectural Equilibrium: A Framework for Robust Predictions," *Theoretical Economics*, 8 (2), 467–501. [570]
FRIEDENBERG, A. (2017): "Bargaining Under Strategic Uncertainty: The Role of Second-Order Optimism," Technical Report, University of Arizona. [570]
FUDENBERG, D., AND Y. KAMADA (2015): "Rationalizable Partition-Confirmed Equilibrium," *Theoretical Economics*, 10 (3), 775–806. [570,585]
GALEOTTI, A., AND S. GOYAL (2010): "The Law of the Few," *American Economic Review*, 100 (4), 1468–1492. [569]
GALEOTTI, A., S. GOYAL, M. JACKSON, F. VEGA-REDONDO, AND L. YARIV (2010): "Network Games," *Review of Economic Studies*, 77 (1), 218–244. [569,573]
GOLUB, B., AND M. JACKSON (2010): "Naïve Learning in Social Networks and the Wisdom of Crowds," *American Economic Journal: Microeconomics*, 2 (1), 112–149. [569]
JACKSON, M. (2008): *Social and Economic Networks*. Princeton University Press. [567]
JACKSON, M., AND A. WOLINSKY (1996): "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 71 (1), 44–74. [569,586]
LOBEL, I., AND E. SADLER (2015): "Information Diffusion in Networks Through Social Learning," *Theoretical Economics*, 10 (3), 807–851. [569]
MARMAROS, D., AND B. SACERDOTE (2002): "Peer and Social Networks in Job Search," *European Economic Review*, 46, 870–879. [569]
MORRIS, S., AND H. S. SHIN (1998): "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 88 (3), 587–597. [570]
OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. MIT Press. [579]
PEARCE, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52 (4), 1029–1050. [570,572,587]
POWELL, J., AND C. THYNE (2011): "Global Instances of Coups From 1950 to 2010: A New Dataset," *Journal of Peace Research*, 48 (2), 249–259. [568]
RUBINSTEIN, A. (1991): "Comments on the Interpretation of Game Theory," *Econometrica*, 59, 909–924. [587]
RUBINSTEIN, A., AND A. WOLINSKY (1994): "Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability," *Games and Economic Behavior*, 6, 299–311. [570,584]
SELTEN, R., AND M. WARGLIEN (2007): "The Emergence of Simple Languages in an Experimental Coordination Game," *Proceedings of the National Academy of Sciences*, 104 (18), 7361–7366. [586]
SHADMEHR, M., AND D. BERNHARDT (2017): "Vanguards in Revolution," Technical Report, University of Calgary & University of Illinois Urbana–Champaigne. [583]
TSAKAS, E. (2013): "Pairwise Epistemic Conditions for Correlated Rationalizability," *Mathematical Social Sciences*, 66 (3), 379–384. [570]
WEBER, R., AND C. CAMERER (2003): "Cultural Conflict and Merger Failure: An Experimental Approach," *Management Science*, 49 (4), 400–415. [586]

Co-editor Joel Sobel handled this manuscript.

Manuscript received 15 August, 2017; final version accepted 9 October, 2018; available online 9 October, 2018.