Strategic diffusion in networks through contagion [*]

Daniel C. Opolot[1] and Théophile T. Azomahou[2]

[1]University of Cape Town
[2]University Clermont Auvergne, CNRS, CERDI

Abstract

This paper studies the diffusion of products and behaviour with coordination effects through social networks when agents are myopic best responders. We develop a new network measure, *the contagion threshold*, that determines when a p-dominant action—an action that is a best response when adopted by at least proportion p of an agent's opponents—spreads to the whole population starting from a group of players whose size is smaller than half and independent of the population size. We show that a p-dominant action spreads to the whole network whenever the contagion threshold of that network is greater or equal to p. We then show that in settings where agents regularly or occasionally experiment and choose non-optimal actions, there exists a threshold level of experimentation below which a p-dominant action is chosen with the highest probability in the long run. This result implies that targeted contagion, a network-wide diffusion of actions initiated by targeting agents, is justified even in settings where agents' decisions are noisy.

Keywords: Contagion, stochastic evolution, networks, experimentation, targeting

JEL Classification: C72, C73, D83, D85

1   Introduction

The adoption of new products and behaviour through social influence is a well-documented phenomenon.[1] This phenomenon provides a justification for *targeted contagion*, a process by which a firm initiates a cascading adoption of a product from a relatively small group (compared to the population size) of consumers, called *initial adopters*, to the whole population. The feasibility of targeted contagion depends on the structure of interactions and payoffs, and the behavioural assumptions that underlie individual decision processes.[2]

This paper studies how the interaction structure and payoffs interactively determine the feasibility of targeted contagion of products and behaviour with coordination effects (i.e. where individuals benefit from coordinating their activities by making the same decisions).[3] The existing literature in this topic focuses on examining contagion in unbounded networks (Morris, 2000;

---

[*] Daniel C. Opolot: `opodanchris@gmail.com`; Théophile T. Azomahou: `theophile.azomahou@uca.fr`.

[1]For example, Bass (1969), Rogers (2003) and Young (2009) document evidence of the diffusion of consumer durables and new products through social influence.

[2]For example, Morris (2000), Alós-Ferrer and Weidenholzer (2008), Galeotti and Goyal (2009), Campbell (2013), Goyal et al (2014), Tsakas (2017) and Beaman et al (2018) all demonstrate how the payoff and interaction structures, and the underling behavioural assumptions determine the feasibility of targeted contagion.

[3]Examples of products and behaviour with coordination effects include information technologies, legal standards, social norms, and political actions such as protests and tactical voting.

Oyama and Takahashi, 2015). However, most social interactions are best described by finite networks, and most of the simplifying assumptions that are applicable to contagion dynamics in unbounded networks do not hold for finite networks.[4] Moreover, the literature does not address the question of whether targeted contagion is economically reasonable in settings where individuals regularly or occasionally experiment with their choices (see for example McKelvey and Palfrey (1995), Anderson et al (2001) and Ellison (2006) for evidence of noisy decision making). In such settings, targeted contagion is economically reasonable if the level of experimentation does not prevent actions from spreading through contagion. To address these gaps in the literature, we examine the following two questions: For a diffusion process on networks where agents are best responders and payoffs exhibit coordination effects, which actions are contagious and when is contagion feasible? In situations where individuals experiment with their choices, when is it economically reasonable to engage in targeted contagion?

To address these questions, we consider two related evolutionary game models. In Model 1, *evolution with best response*, we consider a diffusion process where agents, who are *myopic* best responders, interact through a finite network and revise their actions over time (Morris, 2000; Oyama and Takahashi, 2015). Here, myopia means that agents choose best responses to the most recent actions taken by their network neighbours. It captures the notion of *bounded rationality*; that is, in the real world, individuals have short memories and do not worry about the long-run consequences of their actions (Egidi et al, 1992).

In Model 2, *evolution with best response and mutations*, we consider a diffusion process where agents follow the structural and behavioural assumptions of Model 1, but in addition, they regularly or occasionally make mistakes and choose actions that are not best responses (Young, 1993; Ellison, 2000). Individual mistakes are either due to deliberate experimentation where agents try-out new actions (possibly because of lack of complete information about the game), or due to random errors in action implementation leading agents to choose unintended actions.

We use Model 1 to examine which actions (of a coordination game) are contagious and when contagion is feasible, and Model 2 to examine when targeted contagion is economically reasonable in settings where agents' decisions are noisy. Formally, an action is contagious on a *finite network* if: (*i*) it can spread from a small group of initial adopters, whose size is independent of the population size, to the whole population; (*ii*) it is uninvadable – it can spread from a strictly smaller group (of initial adopters) than is needed to leave a state where it is adopted by the whole population. In Model 2, individual mistakes ensure that every configuration of actions is played with a positive probability at any period. We define the *long-run equilibrium* of this model as the configuration of actions that is played with the highest probability in the long run. Targeted contagion is then said to be *economically reasonable* if the contagious action in Model 1 is also the long-run equilibrium of Model 2 (i.e. if the long-run equilibrium is a configuration containing only the contagious action).

We develop a new network measure – the *contagion threshold* – that determines when con-

_____

[4]For example, when the network is unbounded, an action is said to spread contagiously if it can spread to the whole population starting from a finite group of initial adopters (Morris, 2000). However, when the network structure is finite, it is necessary to have knowledge of the number and identity of agents that can trigger the contagious spread of an action.

tagion is feasible in finite networks. Let a_j be the *p-dominant* action/strategy of a coordination game (i.e. a strategy that is a best response when played by at least proportion p of an agent's neighbours). The contagion threshold of a finite network is the maximum value of p such that action a_j spreads contagiously in that network starting from the *second-neighbourhood* (i.e. the set of all agents within two steps away from a given agent, with that agent included) of any agent. This definition is similar to the definition of the contagion threshold for unbounded networks and 2×2 coordination games in Morris (2000) — the contagion threshold of an unbounded network is the maximum p such that a_j spreads contagiously from a *finite group* of agents. The difference arises because for contagion dynamics in finite networks, it is necessary to have knowledge of the size and identity of the initial adopters that trigger the contagious spread of a p-dominant action.

We show, for Model 1, that a p-dominant action of a coordination game, if it exists, is contagious on any sufficiently large and *strongly connected* network (i.e. a network where every pair of agents is connected through some path) with the contagion threshold greater or equal to p. The set of initial adopters that trigger the contagious spread of a p-dominant action corresponds to the smallest second-neighbourhood of the network. The smallest second-neighbourhood of any network is independent of the population size and it can be as small as three agents. More generally, the smallest second-neighbourhood is smaller for networks with a lower density of connections.[5] This result suggests that the cost of targeted contagion, measured as the number of initial adopters that trigger the contagious spread of a p-dominant action, is lower for sparsely connected than highly connected networks. We discuss this and other implications of these results in more detail in Section 3.

For Model 2, we show that there exists a threshold level of noise below which targeted contagion is economically reasonable. Specifically, below the threshold level of noise, a configuration containing only a p-dominant action is the unique long-run equilibrium in a strongly connected network with the contagion threshold greater or equal to p. Above the threshold level of noise, noisy dynamics becomes more prominent relative to best response dynamics, which makes contagion a less relevant means of spreading an action to the whole population. Targeting therefore becomes economically unreasonable because targeted agents will ultimately switch to other actions through experimentation.

This paper is closely related to the literature that examines how the network structure affects the contagious spread of actions that exhibit coordination effects (Morris, 2000; Alós-Ferrer and Weidenholzer, 2008; Oyama and Takahashi, 2015).[6] Morris (2000) shows that, for best response dynamics on unbounded networks, a p-dominant strategy of a 2×2 coordination game is contagious if p is less or equal to the contagion threshold. Alós-Ferrer and Weidenholzer (2008) consider a model of imitation dynamics and show that, with some restrictions on the information structure, a Pareto-dominant equilibrium of a 2×2 coordination game is contagious on any network that is strongly connected. Oyama and Takahashi (2015) consider a model of best response dynamics on unbounded networks and show that either a risk-dominant or

---

[5]The network density is the number of links connecting agents divided by the total number of links possible.

[6]There is a related literature on diffusion in the presence of coordination effects, for example (López-Pintado et al, 2008), (Jackson and Yariv, 2007), (Sundararajan, 2007), (Galeotti et al, 2010) and (Galeotti and Goyal, 2009). However, these papers study binary choice diffusion processes on random networks.

a Pareto-dominant strategy is contagious in the presence of a third dominated strategy. Like Oyama and Takahashi (2015), we generalize the analysis in Morris (2000) to multiple strategy coordination games. However, Oyama and Takahashi (2015), focus on two types of networks – linear and non-linear networks – so that a strategy is contagious if it can spread contagiously in either of these two networks. Our analysis is richer in that we establish the conditions under which a strategy is contagious on any given network. A more fundamental difference between our analysis and Morris (2000) and Oyama and Takahashi (2015) is that we examine contagion in finite networks and provide steps for computing the smallest number of initial adopters needed to trigger contagion.

Our analysis and results from Model 2 are related to the literature on stochastic evolutionary game dynamics in networks (Ellison, 1993, 2000; Blume, 1995; Berninghaus and Schwalbe, 1996; Young, 1998; Lee and Valentinyi, 2000; Lee et al, 2003; Alós-Ferrer and Weidenholzer, 2007; Peski, 2010; Opolot, 2018, 2020). Except for Peski (2010) and Opolot (2018, 2020), these papers focus on identifying the *stochastically stable states* (i.e. states that are played with a positive probability at the limit of noise) by considering specific network structures, such as a ring and 2-dimensional grid networks. Peski (2010) considers a model of evolution with best response and mutations and shows that a p-dominant strategy is uniquely stochastically stable in networks with the smallest odd degree k_0 (i.e. the size of the smallest odd first-order neighbourhood) satisfying the condition $p \leq \frac{1}{2}(1 - 1/k_0)$. Following the same framework, Opolot (2018) shows that a p-best response set (i.e. a subset of strategies of a coordination game that are best responses when adopted by at least proportion p of an agent's neighbours) is uniquely stochastically stable in any strongly connected network with the contagion threshold greater or equal to p; and Opolot (2020) generalizes these results by establishing the conditions under which the smallest iterated p-best response set is uniquely stochastically stable. In contrast, we establish the conditions under which a p-dominant strategy is not only stochastically stable, but also played with the highest probability in evolutionary models with high noise levels. We then use this result as a justification for targeted contagion in diffusion environments where agents' decisions are not strictly optimal.

The remainder of the paper is organized as follows. In Section 2, we introduce Model 1 and Model 2. Section 3 derives our main results on contagion. In Section 4 we derive the main results of Model 2 and discuss their implications for targeted contagion. Sections 5 and A offer concluding remarks and the implications of our results for the speed of learning respectively. All lengthy proofs are contained in the Appendix.

2 The model

2.1 Network, actions and payoffs

We model the diffusion of strategies (representing actions or products) of a coordination game in social networks using an evolutionary game theory framework. We consider a finite set of players, $N = \{1, \cdots, i, \cdots, n\}$, connected through a social network represented by a graph $G(N, E)$, where E is the set of edges connecting players in N. Let G be an $n \times n$ interaction matrix representing all connections in $G(N, E)$, where $G_{ik} = 1$ if a link exists from i to k, and

zero otherwise. We place the following restrictions on G:

(i) G is undirected and unweighted, that is, either $G_{ik} = G_{ki} = 1$, or $G_{ik} = G_{ki} = 0$.

(i) G is strongly connected, that is, for every pair of agents $i, k \in N$, there exists a path of (undirected) links from i to k.

Both restrictions are simplifications that ensure that the evolutionary process (described below) converges. Directed and/or non-strongly connected networks can lead to cyclic interactions and/or isolated nodes, which inhibits convergence. Let N_i be the set of direct neighbours of i, that is, $N_i = \{k \mid G_{ik} = 1\}$; and let n_i be the cardinality of N_i, also commonly referred in the social networks literature as the *degree* of i.

Every player plays a coordination game with each neighbor. A coordination game consists of a set of strategies, $A = \{a_1, \cdots, a_j, \cdots, a_m\}$, which we assume to be identical for all players, and an $m \times m$ payoff matrix, U. Let $u(a_j, a_l)$ be the payoff of playing strategy a_j against an opponent playing strategy a_l. The double, (A, U), is a coordination game if for each $a_j \in A$, $u(a_j, a_j) \geq u(a_l, a_j)$ for all $a_l \neq a_j$. We focus on *strict symmetric coordination games* where $u(a_j, a_l)$, for all $a_j, a_l \in A$, is identical for all players, and for each $a_j \in A$, $u(a_j, a_j) > u(a_l, a_j)$ for all $a_l \neq a_j$.

Let Σ be the set of all mixed strategies over A so that, for any $\sigma \in \Sigma$, $\sigma(a_j)$ is the mass that σ places on a_j. For player i, we write $\sigma_i = (\sigma_i(a_1), \cdots, \sigma_i(a_m))$ for the distribution over A representing the proportion of i's neighbours playing each pure strategy. We consider linear payoffs where the total payoff that i receives from playing strategy a_j against σ_i is

$$U_i(a_j \mid \sigma_i) = \sum_{a_l \in A} \sigma_i(a_l) u(a_j, a_l). \tag{1}$$

We refer to the quadruple (A, U, N, G) as a *local interaction symmetric coordination game*.

2.2 Behaviour and dynamics

Given the local interaction symmetric coordination game (A, U, N, G), we consider an evolutionary process where players simultaneously and independently revise their strategies at discrete times $t = 1, 2, \cdots$. This evolutionary process suitably models the diffusion of products where agents can replace the old, spoilt or no longer preferred products with new ones of the same type/brand, or with different competing products. We consider two related evolutionary models that capture realistic behavioural assumptions about how people react to observed behaviour:

(i) **Model 1:** *evolution with best response*, where players choose strategies that are best responses to the strategy profile of their direct neighbours.

(ii) **Model 2:** *evolution with best response and mutations (BRM)*, where, in addition to the best response behaviour of Model 1, players make mistakes with a positive probability and choose strategies that are not best responses. Individual mistakes can be thought of as deliberate experimentation with new strategies (i.e. trying-out new strategies possibly due to lack of complete information about the game), or random errors in strategy implementation leading a player to choose an unintended strategy.

As customary in the literature of evolutionary game theory, we assume that players are *myopic* best responders. That is, at every period, t, each player chooses a strategy that is a best response to the distribution of strategies in her neighbourhood at period $t-1$. The assumption of myopia captures the notion that, in the real world, individuals have short memories and are incapable of keeping track of the entire history of play, and that they do not worry about the long run consequences of their strategy choices (Egidi et al, 1992).

To formalize these ideas, let lowercase bold letters e.g. $\mathbf{x}, \mathbf{y}, \mathbf{z}, \cdots$ denote vectors representing profiles (configurations) of strategies. Let x^i (i.e. the ith element of configuration $\mathbf{x}$) be the strategy that player $i \in N$ plays in configuration $\mathbf{x}$. Given the local interaction symmetric coordination game (A, U, N, G), we write $\mathbf{X}$ for the set of all possible strategy configurations. The cardinality of $\mathbf{X}$ is m^n, where m is the number of strategies. Let $\sigma_i(a_l; \mathbf{x})$ be the proportion of i's neighbours playing a_l in configuration $\mathbf{x}$, and let $\sigma_i(\mathbf{x}) = (\sigma_i(a_1; \mathbf{x}), \cdots, \sigma_i(a_m; \mathbf{x}))$ be the distribution over A representing the proportion of i's neighbours playing each strategy under configuration $\mathbf{x}$. Let $\mathbf{x}_t$ be the strategy configuration at period t. Then, following (1), the payoff to i for playing strategy a_j against distribution $\sigma_i(\mathbf{x}_t)$ is $U_i(a_j, \mathbf{x}_t) = \sum_{a_l \in A} \sigma_i(a_l; \mathbf{x}_t) u(a_j, a_l)$. The set of strategies that are best responses to $\sigma_i(\mathbf{x}_t)$ is defined as

$$BR_i(\mathbf{x}_t) = \left\{ a_j : U_i(a_j, \mathbf{x}_t) \geq U_i(a_l, \mathbf{x}_t), \forall a_l \neq a_j \right\}$$

Associated with each $BR_i(\mathbf{x}_t)$ is $BR_i(a_j; \mathbf{x}_t)$, which is the probability that i chooses a_j through best response given configuration $\mathbf{x}_t$. The following assumption places structure on these probabilities.

Assumption 1. *For any $\mathbf{x} \in \mathbf{X}$, if the cardinality of $BR_i(\mathbf{x})$ is greater than one for some $i \in N$, then there exists some tie-breaking rule where $BR_i(a_j; \mathbf{x}) = 1$ for some $a_j \in BR_i(\mathbf{x})$ and $BR_i(a_l; \mathbf{x}) = 0$ for all other $a_l \in BR_i(\mathbf{x})$ and $a_l \neq a_j$.*

Assumption 1 implies that whenever a player is indifferent between two or more strategies, some tie-breaking rule ensures that only one of the strategies is played with probability one. This assumption is a simplification that, although not necessary for Model 1, is used to derive the equilibrium conditions for Model 2.[7]

A tie-breaking rule in the context of Assumption 1 means that, when indifferent between a set of strategies, a player chooses one strategy among this set based on some strategy characteristics. There are several examples of tie-breaking rules that satisfy the conditions in Assumption 1. First, we could let players choose a strategy within $BR_i(\mathbf{x})$ that is also contained in configuration $\mathbf{x}$ and is played by the highest proportion of neighbours: that is, given configuration $\mathbf{x}$ and the respective distribution $\sigma_i(\mathbf{x})$, player i plays a strategy $a_j \in BR_i(\mathbf{x})$ for which $\sigma_i(a_j; \mathbf{x}) > \sigma_i(a_l; \mathbf{x})$ for all $a_l \in BR_i(\mathbf{x}) \backslash a_j$.

Second, we can let players choose some strategy in $BR_i(\mathbf{x})$ that dominates other strategies within $BR_i(\mathbf{x})$ under some criteria. For example, we can choose pairwise domination where $a_j \in BR_i(\mathbf{x})$ is chosen over $a_l \in BR_i(\mathbf{x})$ if its average payoff is higher than for a_l whenever it is played by at least half of neighbours and the rest play a_l: that is, $a_j \in BR_i(\mathbf{x})$ is chosen

---

[7]An alternative consideration is to assume that each strategy in $BR_i(\mathbf{x})$ is played with a uniform probability $\frac{1}{|BR_i(\mathbf{x})|}$, where $|BR_i(\mathbf{x})|$ is the cardinality of $BR_i(\mathbf{x})$. This probability structure would still generate the same results for Model 1, but would lead to less precise analytical results for Model 2.

over $a_l \in BR_i(\mathbf{x})$ if $U_i(a_j, \sigma) > U_i(a_l, \sigma)$ for all σ with $\sigma(a_j) \geq \frac{1}{2}$ and $\sigma(a_l) = 1 - \sigma(a_j)$. Such a relationship between any pair of strategies is called *pairwise risk-dominance* (i.e. a_j is *pairwise risk-dominant* relative to a_l). Thus, a player chooses a_j from $BR_i(\mathbf{x})$ if it is pairwise risk-dominant relative to all other strategies in $BR_i(\mathbf{x})$. Overall, to generate the conditions in Assumption 1, we can apply one or a combination of tie-breaking rules.

Model 1: Given the above definitions and assumptions, the probability that configuration $\mathbf{x}$ is followed by configuration $\mathbf{y}$ at period $t + 1$, denoted by $P(\mathbf{x}, \mathbf{y})$, is given by

$$P(\mathbf{x}, \mathbf{y}) = \prod_{i=1}^{n} BR_i\left(y^i; \mathbf{x}_t = \mathbf{x}\right) \tag{2}$$

where the product on the right hand side of (2) follows because players revise strategies simultaneously and independently.

Model 2: The probability, $\mathbb{P}_i(a_j; \mathbf{x}_t)$, that player i chooses a_j (in period $t + 1$) given configuration $\mathbf{x}_t$, is

$$\mathbb{P}_i(a_j; \mathbf{x}_t) = \frac{1}{m}\exp(-\beta) + (1 - \exp(-\beta))BR_i(a_j; \mathbf{x}_t) \tag{3}$$

The first term on the right hand side of (3), $\frac{1}{m}\exp(-\beta)$, captures individual choices due to mistakes/mutations, and the second term captures best response dynamics. Specifically, a player follows a best response behaviour with probability $(1 - \exp(-\beta))$, and with probability $\exp(-\beta)$, a player mutates and randomly picks any strategy (i.e. with a uniform probability $\frac{1}{m}$).

A closer examination of (3) reveals that as β increases to infinity, players choose best responses with higher probability, and as β tends to zero, agents' choices become more random. The exact value of β may depend on the interaction environment, and may also vary across players. For the former, our analysis examines the effects of varying β on equilibrium behaviour. However, for simplicity, we assume that β is identical across players.

Analogously to (2), the probability, $P_\beta(\mathbf{x}, \mathbf{y})$, that configuration $\mathbf{x}$ is followed by configuration $\mathbf{y}$ in Model 2 is given by

$$P_\beta(\mathbf{x}, \mathbf{y}) = \prod_{i=1}^{n} \mathbb{P}_i\left(y^i; \mathbf{x}_t = \mathbf{x}\right) \tag{4}$$

The dynamics defined by the transition probabilities in (2) and (4) both follow a *stationary Markov chain* on the configuration space $\mathbf{X}$. Let P and P_β denote the respective Markov transition matrices with $P(\mathbf{x}, \mathbf{y})$ and $P_\beta(\mathbf{x}, \mathbf{y})$ as typical elements of P and P_β respectively. We refer to (A, U, N, G, P) and (A, U, N, G, P_β) as the *best response diffusion process* and *best response with mutation diffusion process* corresponding to Model 1 and Model 2 respectively. Given these two models, we seek to:

(i) use Model 1 to establish conditions under which contagion occurs on a given network;

(ii) establish conditions under which targeting is economically reasonable in a framework where agents experiment with their choices.

2.3 Equilibrium behaviour

The equilibrium behaviour of Model 1 is represented by its *limit (absorbing) sets.* A subset $L \subseteq \mathbf{X}$ of states (strategy configurations) is an *absorbing set* of a Markov chain if, once entered, is

never exited. If an absorbing set is a singleton then it is called an *absorbing state* or *conventions*; that is, any state $\mathbf{x} \in \mathbf{X}$ for which $P(\mathbf{x}, \mathbf{x}) = 1$ is an absorbing state. If an absorbing set contains more than one state, then it is referred to as an *absorbing cycle.* For example, a pair of states $\mathbf{x}$ and $\mathbf{y}$ form an absorbing cycle if $P(\mathbf{x}, \mathbf{y}) = 1$ and $P(\mathbf{y}, \mathbf{x}) = 1$.

Let $\mathbf{L}$ be a set of all absorbing sets (i.e. all absorbing states and absorbing cycles) of P. The composition of $\mathbf{L}$ depends on the payoff and network structures. Since (A, U) is a strict symmetric coordination game, $\mathbf{L}$ consists of all *monomorphic configuration* (i.e. where all players coordinate on the same strategy). But depending on the network structure, $\mathbf{L}$ may also contain absorbing cycles and configurations where strategies co-exist. For each $a_l \in A$, we write $\mathbf{a}_l$ for the monomorphic absorbing configuration where all players coordinate on a_l. Associated with each $L \subset \mathbf{L}$ is the basin of attraction, $D(L)$, which is the set of all configurations from which (A, U, N, G, P) converges to L.

For Model 1, the contagion analysis involves examining the stability of monomorphic absorbing configurations (i.e. examining the costs, measured in terms of the number of mutations, of reaching and leaving an absorbing configurations) to determine which configuration contains a contagious strategy. In doing so, we also establish the conditions under which contagion is feasible.

For Model 2, the presence of mutations ensures that every strategy configuration is visited multiple times in the long-run. That is, unlike Model 1 where (A, U, N, G, P) eventually settles in some absorbing set, in Model 2, each configuration is reached with a positive probability at any given time. Thus, a suitable measure of equilibrium behaviour in Model 2 is the *stationary distribution*, denoted by π_β. The stationary distribution of a Markov chain, if it exists, describes the fractional amount of time the process spends in each configuration in the long-run; or equivalently, the probability with which each configuration is visited in the long-run.

Formally, let $\mathbf{q}_0$ be an m^n-row vector, where m^n is the size of the state space $\mathbf{X}$ of (A, U, N, G, P_β), representing an initial distribution of the Markov chain of Model 2. For example, if the chain starts from configuration $\mathbf{x}$, then $\mathbf{q}_0$ is a vector of all zeros except a one in configuration $\mathbf{x}$. Since the transition matrix P_β is homogeneous (i.e. independent of time), the distribution of the Markov chain corresponding to (A, U, N, G, P_β) after t iterations is $\mathbf{q}_t = \mathbf{q}_0 P_\beta^t$. The stationary distribution π_β is then an m^n-vector defined as $\pi_\beta = \lim_{t \to \infty} \mathbf{q}_0 P_\beta^t$. This limit exists and is unique because the Markov chain associated with (A, U, N, G, P_β) is ergodic (i.e. (A, U, N, G, P_β) has a unique absorbing set since every configuration is visited with a positive probability).

Thus, for Model 2, we use the notion of *long-run equilibrium*, which is the set of strategy configurations that are visited most often in the long run by the Markov chain associated with $(A, U, N, G, P_\varepsilon)$. Specifically, a subset $L^* \subset \mathbf{L}$ is the long-run equilibrium of (A, U, N, G, P_β) if $\pi_\beta(L^*) \geq \pi_\beta(L)$ for all $L \subseteq \mathbf{L} \backslash L^*$. Our definition of long-run equilibrium is more general than the notion of *stochastic stability* often used in evolutionary game models. A subset $L^* \subset \mathbf{L}$ is stochastically stable if $\lim_{\beta \to \infty} \pi_\beta(L^*) > 0$, which implies that there exists a threshold value of β above which the long-run equilibria of (A, U, N, G, P_β) are stochastically stable. We show that this threshold value of β also corresponds to the value above which best response dynamics dominates noisy dynamics, so that targeting players is economically reasonable.

3 Diffusion through contagion

We aim to generalize the notion of contagion defined by Morris (2000) for 2×2 coordination games played on unbounded networks to multiple-strategy coordination games played on finite networks. For a 2×2 symmetric coordination game with strategies a_j and a_l played on an unbounded network, strategy a_l is contagious if: (i) starting from a strategy profile where all players play a_j, a_l *spreads contagiously* from a *finite group* of initial adopters; (ii) it is *uninvadable* – once a convention where all players play a_l is established, it should not be possible to leave it with a finite group of deviants. These concepts can be extended to multiple-strategy coordination games and finite networks through the following series of conceptual definitions.

Definition 1. *A sequence of strategy profiles $\{\mathbf{x}_t\}_{t=0}^{\bar{t}}$ of (A, U, N, G, P), for some $\bar{t} \geq 2$, is a best response sequence if it satisfies the following properties: (i) for all $1 \leq t < \bar{t}$, there exists at least one $i \in N$ such that $x_t^i \neq x_{t-1}^i$; (ii) if $x_t^i \neq x_{t-1}^i$, then $x_t^i \in BR(\sigma_i(\mathbf{x}_{t-1}))$.*[8]

Property (i) of Definition 1 requires that at least one player must switch a strategy at each period – this follows because we consider a dynamic process with a simultaneous revision protocol. Property (ii) requires players to switch strategies through best response dynamics.

Definition 2. *Let (A, U, N, G, P) start from some $\mathbf{a} \in \mathbf{L}\backslash\mathbf{a}_l$. Strategy a_l spreads contagiously from a subset of players, $N(\mathbf{a} \rightarrow \mathbf{a}_l) \subset N$, if there exists some $\bar{t} \geq 2$ such that every best response sequence $\{\mathbf{x}_t\}_{t=0}^{\bar{t}}$ with $\mathbf{x}_0 = \mathbf{a}$ and $x_1^i = a_l$ for all $i \in N(\mathbf{a} \rightarrow \mathbf{a}_l)$ satisfies $x_{\bar{t}}^i = a_l$ for all $i \in N$.*

Let $n(\mathbf{a} \rightarrow \mathbf{a}_l)$ be the cardinality of $N(\mathbf{a} \rightarrow \mathbf{a}_l)$. According to Definition 2, when (A, U, N, G, P) starts from $\mathbf{a}$, strategy a_l spreads contagiously from a subset of players, $N(\mathbf{a} \rightarrow \mathbf{a}_l)$, if it spreads to the whole network through best response. If $N(\mathbf{a} \rightarrow \mathbf{a}_l)$ is the smallest set from which a_l can spread contagiously, then we refer to it as the set of *initial adopters* of a_l needed to trigger a network-wide adoption of a_l starting from $\mathbf{a}$.

Definition 3. *A Strategy, a_l, spreads contagiously on network G if for every $\mathbf{a} \in \mathbf{L}\backslash\mathbf{a}_l$, there exists some $N(\mathbf{a} \rightarrow \mathbf{a}_l) \subset N$ with $n(\mathbf{a} \rightarrow \mathbf{a}_l) < \frac{n}{2}$, and independent of n, such that a_l spreads contagiously from $N(\mathbf{a} \rightarrow \mathbf{a}_l)$.*

For unbounded networks considered by Morris (2000) and Oyama and Takahashi (2015), it is sufficient to require $n(\mathbf{a} \rightarrow \mathbf{a}_l)$ to be finite. However, for finite networks, an equivalent requirement is for $n(\mathbf{a} \rightarrow \mathbf{a}_l)$ to be relatively small (i.e. less than half the population size) and independent of the population size. This implies that as n grows, $n(\mathbf{a} \rightarrow \mathbf{a}_l)$ stays finite. One of the objectives of this paper is to define a network measure that determines when a strategy can spread contagiously on a given network. Once such a measure is defined, it is then possible to derive an upper bound for $n(\mathbf{a} \rightarrow \mathbf{a}_l)$, for all $\mathbf{a} \in \mathbf{L}\backslash\mathbf{a}_l$.

Definition 4. *Let $r(\mathbf{a}_l)$ be the number of mutations required to leave convention $\mathbf{a}_l$. Convention $\mathbf{a}_l$ is uninvadable if, for every $\mathbf{a} \in \mathbf{L}\backslash\mathbf{a}_l$, there exists a subset of initial adopters, $N(\mathbf{a} \rightarrow \mathbf{a}_l)$, such that $r(\mathbf{a}_l) > n(\mathbf{a} \rightarrow \mathbf{a}_l)$, and that $r(\mathbf{a}_l)$ is an increasing function of n.*

---

[8]This definition is similar to Oyama and Takahashi (2015, Definition 1) but different in that we consider simultaneous best response dynamics in finite networks.

Definition 4 ensures that once strategy a_l has spread contagiously to the whole network, convention $\mathbf{a}_l$ should not be easily replaced by another convention through a few mutations. The second condition for uninvadability in Definition 4, (i.e. $r(\mathbf{a}_l)$ must be an increasing function of n) ensures that as n grows, it becomes harder (i.e. many mutations are required) to leave convention $\mathbf{a}_l$; and in unbounded networks, $\lim_{n\to\infty} r(\mathbf{a}_l) = \infty$, so that it is not possible to leave $\mathbf{a}_l$ with a finite number of mutations.

Definition 5. *Given a strict symmetric coordination game* (A, U), *strategy* a_l *is contagious in network* $G(N, E)$ *if it spreads contagiously in* $G(N, E)$ *and is uninvadable.*

From Definition 5, we see that strategy a_l is contagious in a finite network if, for all $\mathbf{a} \in \mathbf{L}\backslash\mathbf{a}_l$, $n(\mathbf{a} \to \mathbf{a}_l)$ is independent of the population size, n, and $r(\mathbf{a}_l)$ is an increasing function of n. As the population size becomes infinitely large, we then have an equivalent definition for contagion in unbounded networks where $\lim_{n\to\infty} n(\mathbf{a} \to \mathbf{a}_l) < \infty$ and $\lim_{n\to\infty} r(\mathbf{a}_l) = \infty$.

Given a network G and any subset of players $S \subset N$, let $N_i(S) = N_i \cap S$ be the set of i's direct neighbours within G that belong to subgroup S. Let $n_i(S)$ be the corresponding cardinality of $N_i(S)$ and $\alpha_i(S) = \frac{n_i(S)}{n_i}$ be the proportion of i's neighbours in S.

Now, let (A, U, N, G, P) start from some $\mathbf{a} \in \mathbf{L}\backslash\mathbf{a}_l$. Then from Definitions 1 and 2, strategy a_l spreads contagiously from subgroup $N(\mathbf{a} \to \mathbf{a}_l)$ if the following two conditions hold. First, a_l must be a *p-dominant strategy*, where, following Morris et al (1995), a strategy is p-dominant if it is a best response to all distributions that place on it a mass of at least p. That is, for a strict symmetric coordination game (A, U), strategy a_l is (strictly) p-dominant if for all $a_k \in A$, and all $\sigma \in \Sigma$ with $\sigma(a_l) \geq p$,

$$\sum_{a_k \in A} \sigma(a_k) u(a_l, a_k) > \sum_{a_k \in A} \sigma(a_k) u(a_j, a_k) \quad \text{for all } a_j \neq a_l$$

In the context of the local interaction symmetric coordination game, a strategy is p-dominant if it is a best response whenever it is played by at least proportion p, or a total of at least $\lceil pn_i \rceil$, of a player's neighbours, where $\lceil x \rceil$ is the least integer greater than or equal to x.

Second, there must exist a sequence of sets of players, $S_1, S_2, \cdots, S_J$, with $S_1 = N(\mathbf{a} \to \mathbf{a}_l)$ and $\bigcup_{j=1}^{J} S_j = N$, whereby for each $i \in S_j$, and all $j = 2, 3, \cdots, J$, $\alpha_i(S_{j-1}) \geq p$. That is, for each player $i \in S_j$, the proportion of i's neighbours that belong to S_{j-1} is greater than p. This inequality ensures that once all players in $S_1 = N(\mathbf{a} \to \mathbf{a}_l)$ play strategy a_l from period $t = 1$ onward, then at $t = 2$, a_l is a best response to all $i \in S_2$, so that all players in $S_1 \cup S_2$ play a_l; at $t = 3$, a_l is a best response to all $i \in S_3$, so that all players in $S_1 \cup S_2 \cup S_3$ play a_l; and so on, until $t = J$ at which all players play strategy a_l.

|	a_1	a_2
a_1	$1, 1$	$0, 0$
a_2	$0, 0$	β, β

Figure 1: A 2×2 strict coordination game where a_2 is a p-dominant strategy, for $p > \frac{1}{1+\beta}$. That is, a_2 is a best response when played by more than proportion $\frac{1}{1+\beta}$ of neighbours.

Example 1. *Consider the 2×2 coordination game in Figure 1 with strategy a_2 as a p-dominant strategy, for $p > \frac{1}{1+\beta}$. First consider a scenario where $\beta = 1.01$, so that a_2 is a $\frac{1}{2}$-dominant strategy. When this game is played on network G_2 of Figure 2a, strategy a_2 is a best response whenever it is played by at least $\lceil \frac{1}{2} \times 2 \rceil = 1$ neighbour. If (A, U, N, G, P) starts from a monomorphic configuration $\mathbf{a}_1$, strategy a_2 will spread contagiously from any pair of two adjacently placed players. This is because, for any pair of adjacently players, say $N(\mathbf{a}_1 \to \mathbf{a}_2) = \{1, 2\}$, there exists a sequence of sets of players, $S_1 = \{1, 2\}, S_2 = \{3, 12\}, S_3 = \{4, 11\} \cdots, S_6 = \{7, 8\}$, whereby for each $i \in S_j$, $\alpha_i(S_{j-1}) = \frac{1}{2} = p$. That is, if all players in S_1 play a_2 from period $t = 1$ onward, then at $t = 2$, all players in $S_1 \cup S_2$ play a_2; at $t = 3$, all players in $S_1 \cup S_2 \cup S_3$ play a_2; and so on, until $t = 6$ at which all players play a_2.*

However, for network G_3 of Figure 2b, when $\beta = 1.01$, strategy a_2 is a best response only when played by at least $\lceil \frac{1}{2} \times 3 \rceil = 2$ neighbours. This implies that for any sequence, $S_1, S_2, \cdots, S_J$, of players in G_3 where the size of S_1 is less than $\frac{n}{2}$, there exists at least one i in some S_j with $\alpha_i(S_{j-1}) = \frac{1}{3} < p = \frac{1}{2}$.[9] Thus, $n(\mathbf{a}_1 \to \mathbf{a}_2) \geq \frac{n}{2}$ and increases with the population size, and according to Definitions 2 and 3, strategy a_2 does not spread contagiously on G_3. Strategy a_2 spreads contagiously on G_3 only when $\beta \geq 2.01$ so that a_2 is $\frac{1}{3}$-dominant.
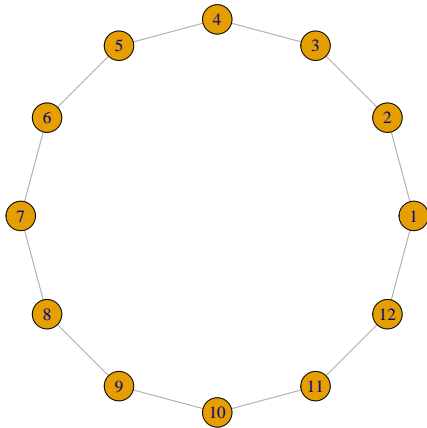
(a) A regular cyclic network, G_2, of degree 2.
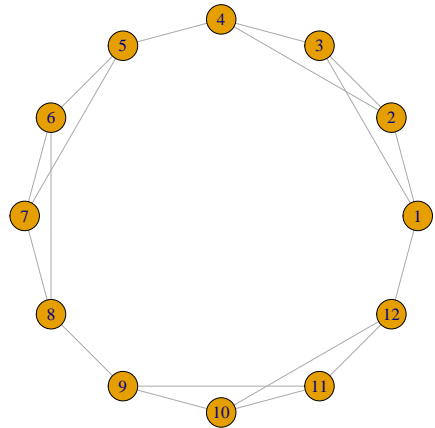
(b) A regular cyclic network, G_3, of degree 3.

Figure 2: Examples of regular cyclic networks

We define a new network measure – the *contagion threshold* – for finite networks that captures the second condition above. This measure not only determines when a *p*-dominant strategy can spread contagiously, but also when it is uninvadable. For each $i \in N$, let B_{i_r} be the *r*th *neighbourhood* of i, the set of all players within distance r (i.e. r steps) from i, with i included, and N_{i_r} be the *r*th-order neighbours of i (i.e. all players at distance r from i). We write b_{i_r} and n_{i_r} for the respective cardinalities of B_{i_r} and N_{i_r}. We then define and compute the contagion threshold of a finite network as follows.

Definition 6. *Given $G(N, E)$:*

---

[9]For example, if $S_1 = \{1, 2, 3, 4\}$, so that $S_2 = \{5, 12\}, S_3 = \{6, 11\}, S_4 = \{7, 10\}$ and $S_5 = \{8, 9\}$, then for players 5 and 12, $\alpha_5(S_1) = \alpha_{12}(S_1) = \frac{1}{3}$. Thus, even if all players in S_1 mutate to a_2 at $t = 1$, a_2 will not be a best response to players 5 and 12 since a_2 is a best response only when it is adopted by at least half of the neighbours.

(i) pick any $i \in N$ and the corresponding B_{i_1};

(ii) for each $r \geq 2$ and $j \in N_{i_r}$, compute $\alpha_j(B_{i_{r-1}})$ and $\alpha_{i_r}^* = \min\limits_{j \in N_{i_r}} \alpha_j(B_{i_{r-1}})$;

(ii) given $i \in N$, compute $\alpha_i^* = \min\limits_{r \geq 2} \alpha_{i_r}^*$.

The contagion threshold of $G(N, E)$, $\eta(G)$, is given by $\eta(G) = \min\limits_{i \in N} \alpha_i^*$.

The definition of the contagion threshold builds on the notion of sequences of sets of players discussed above. For each i with a corresponding first neighbourhood B_{i_1}, Definition 6 requires that along a sequence of the r-order neighbours of i, $N_{i_2}, N_{i_3}, \cdots, N_{i_{d_i}}$, where d_i is the shortest distance from i to the furthest player, $\alpha_j(B_{i_{r-1}}) \geq \eta(G)$ for every $j \in N_{i_r}$ and all $r = 2, 3, \cdots, d_i$. For the two networks in Figure 2, the contagion thresholds are $\eta(G_2) = \frac{1}{2}$ and $\eta(G_3) = \frac{1}{3}$.

Following the above discussion and Example 1, when $p \leq \eta(G)$, a p-dominant strategy can spread contagiously from the set B_{i_2} of any i. Consider the game of Figure 1, and let $\beta = 1.01$ so that a_2 is a $\frac{1}{2}$-dominant strategy. Consider an evolutionary process (A, U, N, G, P) on a strongly connected network G with $\eta(G) = p = \frac{1}{2}$. Then, starting from $\mathbf{a}_1$, strategy a_2 will spread contagiously from $N(\mathbf{a}_1 \rightarrow \mathbf{a}_2) = B_{i_2}$ of any $i \in N$. This is because along the sequence $N_{i_3}, N_{i_4}, \cdots, N_{i_{d_i}}$, $\alpha_j(B_{i_{r-1}}) \geq \eta(G) = \frac{1}{2} = p$ for every $j \in N_{i_r}$ and all $r = 3, 4, \cdots, d_i$, so that once all players in B_{i_2} play a_2 at some period t, all players in N_{i_3} will switch to a_2 at $t + 1$ (because each has at least proportion p of neighbours play a_2 at t); followed by all players in N_{i_4} at $t + 2$; and so on, until the entire network eventually switches to a_2.

Our definition of the contagion threshold for finite networks is closely related to the definition of the contagion threshold for unbounded networks according to Morris (2000). For a 2×2 symmetric coordination game with strategies a_1 and a_2, let strategy a_2 be a p-dominant strategy. Morris (2000) defines the contagion threshold of an unbounded network as the maximum p such that a_2 can spread contagiously in that network. This definition ensures that the contagious spread of a_2 can be triggered from *some finite group* of players. The similarity with Definition 6 is that $\eta(G)$ is the maximum p above which a p-dominant strategy of an $m \times m$ symmetric coordination game can spread contagiously in a finite G.

However, the definition of the contagion threshold for unbounded networks in Morris (2000) does not directly carry-on to finite networks. This is because unlike unbounded networks where it is sufficient to know that contagion can be triggered from some finite group of players, for finite networks, it is necessary to have knowledge of the size and the identity of the smallest group of players from which contagion can be triggered. From Definition 6 and the discussion that follows, the contagious spread of a p-dominant strategy can be triggered from within the smallest second neighbourhood of the network. We further discuss the bounds for the size of the smallest group that triggers the contagious spread of a p-dominant strategy below. But first, the following theorem establishes the conditions for a p-dominant strategy to be contagious on a given network.

Theorem 1. *Given the diffusion process (A, U, N, G, P) on a strongly connected network G, if $a^* \in A$ is a p-dominant strategy of (A, U), then it is contagious in G if $p \leq \eta(G)$ and $b_2^* \leq n^{\frac{3}{5}}$, where $b_2^* = \min_{i \in N} b_{i_2}$. For all $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$, $n(\mathbf{a} \rightarrow \mathbf{a}^*) \leq b_2^* \leq n^{\frac{3}{5}}$.*

Proof. See Appendix B

$\square$

The proof of Theorem 1 follows in two steps. Let $\mathbf{a}^*$ be a monomorphic absorbing state where all players coordinate on strategy a^*. We first show that if a^* is a p-dominant strategy of (A, U), and $p \leq \eta(G)$, then starting from any state $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$, the size of the subgroup from which a^* spreads contagiously in G is bounded from above by b_2^*.

Let (A, U, N, G, P) start from some $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$, and pick any $i \in N$. If all players in B_{i_2} mutate to a^* at $t = 1$, then from $t = 2$ onward, a^* is a best response to all players in $B_{i_2} \cup N_{i_3} = B_{i_3}$ since $\alpha_j(B_{i_2}) \geq \eta(G) \geq p$ for each $j \in B_{i_3}$. That is, since each $j \in B_{i_3}$ has at least proportion p of their interactions with players in B_{i_2}, all of whom play a^* at $t = 1$, and a^* is a best response whenever it is played by at least proportion p of neighbours, it follows that a^* is a best response for each $j \in B_{i_3}$. From $t = 3$ onward, a^* is a best response to all $B_{i_3} \cup N_{i_4} = B_{i_4}$ since $\alpha_j(B_{i_3}) \geq \eta(G) \geq p$ for each $j \in B_{i_4}$; and so on, until $t = d_i - 1$ when the entire network eventually plays a^*.

Now, let B_2^* be the second neighbourhood of player i for whom $b_{i_2} = b_2^*$ (i.e. the cardinality of B_2^* is b_2^*), and let d^* be the shortest distance from the player for whom $b_{i_2} = b_2^*$ to any other player (i.e. $d^* = d_{i^*}$, where i^* is some player for whom $b_{i_2^*} = b_2^*$). Then the above iterative process implies that, for every $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$, there exists a best response sequence $\{\mathbf{x}_t\}_{t=0}^{d^*-1}$ with $\mathbf{x}_0 = \mathbf{a}$ and $x_1^j = a^*$ for all $j \in B_2^*$ satisfying $x_{d^*-1}^j = a^*$ for all $j \in N$. Thus, a^* spreads contagiously on any strongly connected network G whenever $p \leq \eta(G)$. The size of the subgroup from which a^* spreads contagiously is bounded from above by b_2^*.

Note that b_2^* is an upper bound for the minimum number of mutations needed for a^* to spread contagiously on a given network because the exact number of mutations can be much smaller than b_2^*. The exact number of initial adopters of a^* needed to trigger its contagious spread can be computed in two steps:

(i) identify B_2^* by computing B_{i_2} for each $i \in N$ and then picking the one with the smallest cardinality;

(ii) within B_2^*, identify the smallest set of players that should play a^* so that all players in B_2^* eventually switch to play a^* through best response. The cardinality of this subset of B_2^* is then the number of initial adopters.

The second step of the proof of Theorem 1 shows that if a^* is the p-dominant strategy, and $p \leq \eta(G)$, then the number of mutations needed to leave (the basin of attraction of) convention $\mathbf{a}^*$ is greater than $n^{\frac{3}{5}}$. The intuition behind this result is that since each $j \in N_{i_r}$ has $\alpha_j(B_{i_{r-1}}) \geq \eta(G) \geq p$ for all $r \geq 2$, and that $p \leq \frac{1}{2}$, players within a given B_{i_2} will switch to a strategy different from a^* if each has more than proportion $(1 - p) > \frac{1}{2}$ of neighbours playing a strategy different from a^*. Thus, to leave convention $\mathbf{a}^*$, the set of mutants, $R(\mathbf{a}^*)$, must be selected in such a way that, for each $i \in N$, each $j \in B_{i_2}$ has more than half of their neighbours in $R(\mathbf{a}^*)$. This implies that the identification of $R(\mathbf{a}^*)$ is equivalent to the graph theory problem of identifying monopolies (i.e. sets of vertices of a graph containing at least half of the direct and/or indirect interactions of every player). Using the well established results in graph theory

(e.g. Bermond et al (1996, Proposition 4)), it follows that the cardinality of $R(\mathbf{a}^*)$ is greater than $n^{\frac{3}{5}}$.

Since the number of mutations needed to leave convention $\mathbf{a}^*$ is a function of n, it follows that $\mathbf{a}^*$ is uninvadable whenever $b_2^* \leq n^{\frac{3}{5}}$. Thus, a p-dominant strategy, a^*, is contagious on a strongly connected network G whenever $p \leq \eta(G)$ and $b_2^* \leq n^{\frac{3}{5}}$. The following example illustrates a direct application of Theorem 1 to two networks with different contagion thresholds.

Example 2. *Consider the network structures in Figure 3. Network G_a of Figure 3a has a contagion threshold of $\eta(G_a) = \frac{2}{7}$.*[10] *According to Theorem 1, a p-dominant strategy a^* spreads contagiously in G_a if $p \leq \frac{2}{7}$. To identify the minimum number of mutations needed to trigger the contagious spread of a^* in this network, we follow the two steps outlined above. First, we identify the B_2^* sets. There are six B_2^* sets in G_a and we pick the one centred around player 1, $B_2^* = B_{1_2} = \{1, 2, 7, 8, 15, 16, 17, 18\}$.*

Second, we identify the subset of players of B_2^* that must mutate to a^* so that all players in B_2^* eventually play a^* through best response. Consider the case where $p = \frac{2}{7}$ so that, for players $\{1, 2\}$, a^* is a best response (BR) when played by at least $\lceil \frac{2}{7} \times 2 \rceil = 1$ neighbour; for player 7, a^* is a BR when played by at least $\lceil \frac{2}{7} \times 5 \rceil = 3$ neighbours; and for players $\{8, 15, 16, 17, 18\}$, a^* is a BR when played by at least $\lceil \frac{2}{7} \times 6 \rceil = 2$ neighbours. For this scenario, if players $\{1, 2, 7, 8\}$ simultaneously mutate to a^*, then all players in B_{1_2} will eventually switch to a^* through best response. Any other combinations of less than four mutations by players within B_{1_2} may lead to an absorbing cycle containing a^* and other strategies but not convention $\mathbf{a}^*$. Thus, the number of initial adopters that trigger the contagious spread of a^* is four, a small number that is independent of the population size.

For network G_b of Figure 3b, the contagion threshold is $\eta(G_b) = \frac{1}{6}$. Pick $B_2^* = B_{1_2} = \{1, 7, 13, 14\}$. For each $j \in B_{1_2}$, a^* is a BR when played by at least one neighbour. Thus, all players in B_{1_2} will switch to a^* within three steps of iteration once one player, player $\{7\}$, switches to a^*. The corresponding number of initial adopters is thus one, which is a very small number compared to the population size.

For both networks, $\mathbf{a}^*$ is uninvadable since the cardinality of B_2^* is less than $n^{\frac{3}{5}}$. Specifically, for $p \leq \frac{1}{2}$, $R(\mathbf{a}^*)$ must at the least consist of $R(\mathbf{a}^*) = \{7, 8, 9, 10, 11, 12, 13, 14, 23, 24, 25, 26\}$ for network G_a, and $R(\mathbf{a}^*) = \{13, 14, 15, 16, 17, 18, 19, 20, 29, 30, 31, 32\}$ for network G_b. Thus, $r(\mathbf{a}^*) \geq 12$ for both networks.

The above example helps to highlight two interactive aspects of our results which a firm/planner aiming to diffuse a product/behaviour must consider: the contagion threshold versus the number of initial adopters needed to trigger the contagious spread of a strategy. Consider a new firm contemplating to enter a market (or an existing firm aiming to introduce a new product) where two or more products (exhibiting coordination effects) already exist. To diffuse her product through contagion, the firm would first determine the contagion threshold of the network of consumers. The contagion threshold in turn determines the value of p (i.e. the extent to which

---

[10]That is, for each i and corresponding B_{i_1}, each $j \in N_{i_r}$ for all $r \geq 2$ has $\alpha_j(B_{i_{r-1}}) \geq \frac{2}{7}$. For example, if we pick player 1 and corresponding $B_{1_1} = \{7, 8\}$, we see that each $j \in N_{1_2} = \{2, 15, 16, 17, 18\}$ has $\alpha_j(B_{1_1}) \geq \frac{1}{3}$; each $j \in N_{1_3} = \{9, 10, 23, 24\}$ has $\alpha_j(B_{1_2}) \geq \frac{1}{2}$; each $j \in N_{1_4} = \{3, 27, 28, 29, 30\}$ has $\alpha_j(B_{1_3}) \geq \frac{2}{7}$; each $j \in N_{1_5} = \{25, 26\}$ has $\alpha_j(B_{1_4}) = \frac{1}{2}$; each $j \in N_{1_6} = \{19, 29, 21, 22\}$ has $\alpha_j(B_{1_5}) \geq \frac{1}{3}$; each $j \in N_{1_7} = \{11, 12, 13, 14\}$ has $\alpha_j(B_{1_6}) \geq \frac{2}{3}$; and each $j \in N_{1_8} = \{4, 5, 6, 7\}$ has $\alpha_j(B_{1_7}) = 1$.
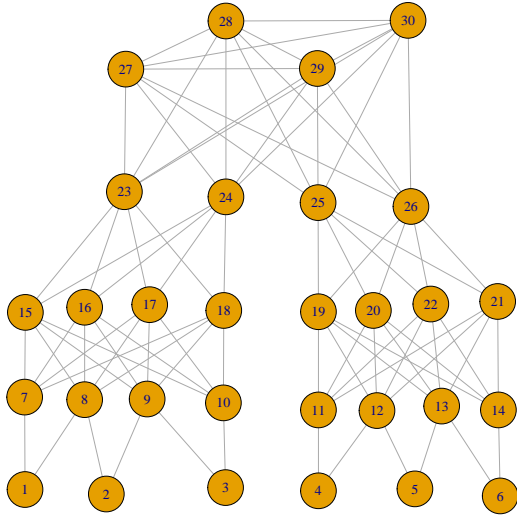
(a) Network G_a with contagion threshold of $\eta(G_a) = \frac{2}{7}$.
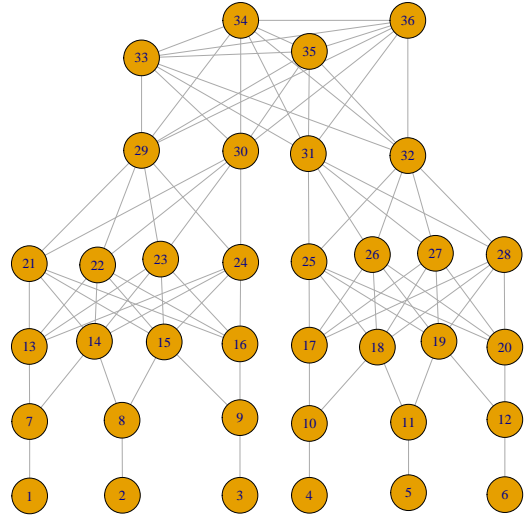
(b) Network G_b with contagion threshold of $\eta(G_b) = \frac{1}{6}$.

Figure 3: Network contagion thresholds.

a new product should dominate the existing alternative products) that ensures that the product can spread contagiously. If the contagion threshold of the network is much smaller than $\frac{1}{2}$, then the firm must incur large initial costs to ensure that p is sufficiently smaller (i.e. the new product must be highly beneficial compared to existing products) than the contagion threshold. The upside is that the smaller p, the smaller the number of initial adopters needed to trigger the contagious spread of a p-dominant strategy, and hence, a firm need not incur large costs on targeting. Conversely, if the contagion threshold is close to $\frac{1}{2}$, then p can also be close to $\frac{1}{2}$ so that the new product only needs to be slightly better than the existing product for it to spread contagiously. The downside is that the firm may have to invest more resources on targeting initial adopters. A firm's objective is thus to evaluate these two types of costs.

Overall, low-density networks (i.e. networks with sparse connection) have lower contagion thresholds. This is because the smallest possible value for the contagion threshold of any strongly connected network G is $\frac{1}{\bar{n}(G)}$, where $\bar{n}(G)$ is the size of the smallest first-order neighbourhood in G. Moreover, low-density networks also have smaller b_2^*, and hence, a small number of initial adopters needed to trigger contagion. This implies that it is relatively less costly to diffuse new products and behaviour to low-density networks. One possible evidence of this implication of our results is in the observed discrepancy between the dynamics in the market for messaging and chat apps (e.g. AOL Instant Messenger, Google talk, VOIP, skype, Kik., whatsapp, snapchat, HipChat and slack), where there is a high turnover and entry rate, versus social networking apps (e.g. Friendster, Myspace and Facebook), where turnover and entry rate is relatively low. Specifically, messaging and chat apps consist of low-density networks where individual interact with (speak, chat or message to) a few others (family and close friends). Social network apps on the other hand consist of large number of connections where some individuals may have as many as tens of thousands of friends. Thus, to invade a network of messaging and chat apps, a new product need not be much more beneficial than those existing in the market. Invading a

Facebook-type network, however, requires a new product to be far better, or simply to offer a different but related services.

Like Oyama and Takahashi (2015), our results extend Morris (2000) who examines contagion for 2×2 coordination games to multiple strategy coordination games. However, in Oyama and Takahashi (2015), a strategy is contagious if it is contagious in *some* unbounded network. As such, it is sufficient to check whether a strategy spreads contagiously in any of two types of networks, linear and non-linear networks. Our analysis is richer in that we establish conditions under which a strategy (i.e. a *p*-dominant strategy) is contagious on any given network. In doing so, we also provide steps for computing the smallest number of initial adopters needed to trigger contagion.

A more fundamental difference between our analysis and Morris (2000) and Oyama and Takahashi (2015) is that we study contagion in finite networks. In particular, Morris (2000) provides bounds for the contagion threshold in unbounded networks but computing the exact contagion threshold of an unbounded network is not a straightforward matter. Ideally, one would select some finite group of players in some region of the network and then iterate over the entire network following similar steps in Definition 6 above. The problem with using this method when the network is unbounded is that there is no obvious means of knowing how many iterations are sufficient for computing the contagion threshold. One would have to assume some form of uniformity across the unbounded network, and hence, the contagion threshold can feasibly be determined for only unbounded regular or close to regular networks (i.e. networks where all players have the same neighbourhood size).

Finally, Alós-Ferrer and Weidenholzer (2008) examine contagion in both bounded and unbounded network but focuses on establishing conditions under which a Pareto-dominant strategy spreads under imitation dynamics – where players play strategies that earned the highest payoff in their neighbourhood in the previous period. Imitation dynamics does not, in most cases, generate intermediate absorbing states where strategies co-exist as best response dynamics does. Thus, the network structure does not affect contagion under imitation dynamics in the same way it does under best response dynamics. Consequently, the network measures that affect contagion under best response dynamics, such as contagion thresholds, are not necessarily relevant for contagion under imitation dynamics.

4   Contagion and best response dynamics with mutations

In Section 3, we established the conditions for contagion in a framework where players strictly follow best response dynamics. The question we ask in this section is the following: is contagion relevant in a framework of noisy best response dynamics? There is plenty of empirical evidence showing that individual decision processes are best described by probabilistic models (McKelvey and Palfrey, 1995; Anderson et al, 2001; Ellison, 2006).

Naturally, the level of noise in individual decisions may vary across settings and it is possible that, in some settings, the amount of noise affects the relevance of the process of contagion. That is, if players frequently make mistakes and play strategies that are not best responses, then the evolutionary process will frequently deviate from paths corresponding to contagion dynamics (i.e. the evolutionary process frequently deviates from best response sequences defined in Definition

1). This may in turn imply that contagious strategies are not played with the highest probability in the long-run. In such a scenario, it will not be economically reasonable to target specific players to aid network-wide diffusion.

To examine the above question, we establish conditions under which a convention corresponding to a contagious strategy is a long-run equilibrium in Model 2. We show that, indeed, under Model 2, there exists a threshold value, $\beta^* > 0$, of the noise parameter β, above which the convention corresponding to the p-dominant strategy is the long-run equilibrium. Above β^*, best response dynamics, and hence, contagion, dominates experimentation and targeting is economically reasonable; conversely, below β^*, players' choices are too noisy and targeting is not economically reasonable. The following definition and notations are used to state the main results of this section.

Definition 7. *Let $g \subset \mathbf{X} \times \mathbf{X}$ be any oriented graph defined within the configuration space $\mathbf{X}$. Then for a subset $W \subset \mathbf{X}$ and its complement $\bar{W}$, we denote by $\Gamma(W)$ a set of all oriented graphs satisfying two conditions: (i) no arrows start from W and exactly one arrow starts from each configuration outside of W, (ii) each $g \in \Gamma(W)$ has no loops.*

From Definition 7, if W is a singleton set, say $W = \{\mathbf{x}\}$, then $\Gamma(\{\mathbf{x}\})$ is a set of all spanning trees of $\mathbf{x}$ (i.e. $\mathbf{x}$-trees). Consider an example where $\mathbf{X} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e}, \mathbf{f}, \mathbf{g}, \mathbf{h}\}$; Figure 4 presents two examples of $\mathbf{g}$-trees.
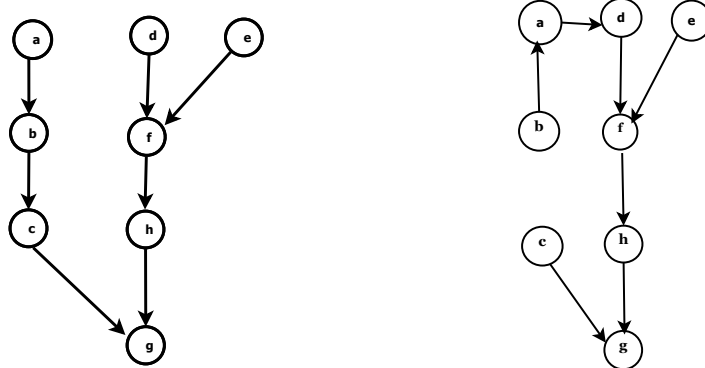

Figure 4: Examples of $\mathbf{g}$-trees, that is $\Gamma(\{\mathbf{g}\})$ graphs and in which the configuration space $\mathbf{X} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e}, \mathbf{f}, \mathbf{g}, \mathbf{h}\}$.

From the definition of $\Gamma(\mathbf{x})$ graphs, every $g \in \Gamma(\mathbf{x})$ spans the entire state space, except for $\mathbf{x}$. Each $\mathbf{y} \in \mathbf{X}$ has only one arrow emanating from it so that each $g \in \Gamma(\mathbf{x})$ has a total of $m^n - 1$ directed edges, where m^n is the cardinality of $\mathbf{X}$. Thus, if we let $\gamma(\mathbf{x}) = \#\Gamma(\mathbf{x})$ be the cardinality of $\Gamma(\mathbf{x})$, then $\gamma(\mathbf{x})$ is identical for any pair of configurations $\mathbf{x}, \mathbf{y} \in \mathbf{X}$; that is, $\gamma(\mathbf{x}) = \gamma(\mathbf{y}) = \gamma$.

The cardinality of $\Gamma(\mathbf{x})$, γ, is a multiple of, and not an exponential function of the cardinality of the state space, m^n. For example, when the size of the state space is 2, $\gamma = 1$; when it is 3, $\gamma = 3$; when it is 4, $\gamma = 15$;[11] and when it is 5, $\gamma = 51$. Thus, the natural logarithm of γ, $\ln(\gamma)$, is a linear function of n.

---

[11]That is, if the state space is $\mathbf{X} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$, then the list of $\Gamma(\mathbf{d})$ graphs is: $\{\mathbf{c} \to \mathbf{b} \to \mathbf{a} \to \mathbf{d}\}, \{\mathbf{a} \to \mathbf{c} \to \mathbf{b} \to \mathbf{d}\}, \{\mathbf{a} \to \mathbf{b} \to \mathbf{c} \to \mathbf{d}\}, \{\mathbf{b} \to \mathbf{a} \to \mathbf{c} \to \mathbf{d}\}, \{\mathbf{b} \to \mathbf{c} \to \mathbf{a} \to \mathbf{d}\}, \{\mathbf{c} \to \mathbf{a} \to \mathbf{d}, \mathbf{b} \to \mathbf{d}\}, \{\mathbf{a} \to \mathbf{c} \to \mathbf{d}, \mathbf{b} \to \mathbf{d}\}, \{\mathbf{c} \to \mathbf{b} \to \mathbf{d}, \mathbf{a} \to \mathbf{d}\}, \{\mathbf{b} \to \mathbf{c} \to \mathbf{d}, \mathbf{a} \to \mathbf{d}\}, \{\mathbf{a} \to \mathbf{b} \to \mathbf{d}, \mathbf{c} \to \mathbf{d}\}, \{\mathbf{a} \to \mathbf{b} \to \mathbf{d}, \mathbf{c} \to \mathbf{d}\}, \{\mathbf{a} \to \mathbf{c}, \mathbf{b} \to \mathbf{c}, \mathbf{c} \to \mathbf{d}\}, \{\mathbf{c} \to \mathbf{a}, \mathbf{b} \to \mathbf{a}, \mathbf{a} \to \mathbf{d}\}, \{\mathbf{a} \to \mathbf{b}, \mathbf{c} \to \mathbf{b}, \mathbf{b} \to \mathbf{d}\}, \{\mathbf{c} \to \mathbf{d}, \mathbf{a} \to \mathbf{d}, \mathbf{b} \to \mathbf{d}\}$

Proposition 1. *For a diffusion model (A, U, N, G, P_β), let Assumption 1 hold. If a^* is contagious on network G, then there exists some $\beta^* \in \left(0, \ln \gamma / (n^{\frac{3}{5}} - b_2^*)\right]$ such that for all $\beta \geq \beta^*$, $\pi_\beta(\mathbf{a}^*) > \pi_\beta(\mathbf{a})$ for all $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$.*

Proof. See Appendix C

$\square$

The proof of Proposition 1 follows in two steps. First, we characterise the structure of stationary distributions using a graph theoretic method of Freidlin and Wentzell (1984). This method represents the stationary distribution of any configuration $\mathbf{x}$ in terms of the probabilities of $\Gamma(\mathbf{x})$ graphs. Let $c(\mathbf{y}, \mathbf{z})$ be the number of mutations involved in the direct transition from $\mathbf{y}$ to $\mathbf{z}$. The total cost $\phi(\mathbf{x}; g)$ of some $g \in \Gamma(\mathbf{x})$ is the sum of costs of all transitions in g. That is, $\phi(\mathbf{x}; g) = \sum_{(\mathbf{y}, \mathbf{z}) \in g} c(\mathbf{y}, \mathbf{z})$. This quantity captures the total cost of reaching $\mathbf{x}$ from every other configuration through paths of graph g.

These costs are directly related to the transition probabilities. The probability that $c(\mathbf{y}, \mathbf{z})$ players simultaneously mutate to play strategies that are not best responses to $\mathbf{y}$ is $\left(e^{-\beta}/m\right)^{c(\mathbf{y}, \mathbf{z})}$. Thus, if the probability of mistakes is relatively small, then the larger $c(\mathbf{y}, \mathbf{z})$ the smaller the probability of the transition from $\mathbf{y}$ to $\mathbf{z}$. The same argument extends to any $g \in \Gamma(\mathbf{x})$: if the probability of mistakes is small, then the higher the cost $\phi(\mathbf{x}; g)$ associated with g, the smaller the probability of reaching $\mathbf{x}$ through paths in g. Using the relationship between stationary distributions and the probabilities of $\Gamma(\mathbf{x})$ graphs provided by Freidlin and Wentzell (1984, Lemma 3.1), we derive bounds for stationary distributions as functions of costs of $\Gamma(\mathbf{x})$ graphs and model parameters.

Second, we show that under Assumption 1, there exists some $\beta^* \in \left(0, \ln \gamma / (n^{\frac{3}{5}} - b_2^*)\right]$ such that for all $\beta \geq \beta^*$, the long-run equilibrium is the strategy configuration with the minimal cost $\Gamma(\mathbf{x})$ graph: that is, for all $\beta \geq \beta^*$, $\mathbf{x} \in \arg \min_{\mathbf{x} \in \mathbf{X}, g \in \Gamma(\mathbf{x})} \phi(\mathbf{x}; g)$ is the long-run equilibrium; and that convention $\mathbf{a}^*$, corresponding to a contagious strategy a^*, has the minimal cost graph.

As discussed above, the intuition behind Proposition 1 is that when the level of experimentation is very high (i.e. β is small and between 0 and β^*), the evolutionary process frequently deviates from the contagion paths (i.e. paths of best response sequences) so that the convention containing only the contagious strategy need not be the long-run equilibrium. For example, when $\beta = 0$, every strategy configuration is visited with equal probability. The upper bound for β^*, $\ln \gamma / (n^{\frac{3}{5}} - b_2^*)$, increases with n. Specifically, $\ln \gamma / (n^{\frac{3}{5}} - b_2^*)$ is of order $(O)(n^{2/5})$ because $\ln \gamma$ is a linear function of n.

The findings of Proposition 1 suggest that targeting is indeed economically reasonable even in settings where players make mistakes and play strategies that are not best responses. More broadly, Proposition 1 suggests that the predictions in evolutionary game theory that employ stochastic stability as a solution concept are admissible for relatively high levels of experimentation. Recall that a configuration $\mathbf{x}$ is stochastically stable if $\lim_{\beta \to \infty} \pi_\beta(\mathbf{x}) > 0$ (Foster and Young, 1990; Kandori et al, 1993; Young, 1993). Proposition 1 states that the level of experimentation need not be very small (i.e. β need not be asymptotically large); it is sufficient that $\beta \geq \beta^*$. And in evolutionary models with a finite and small population size, the level of experimentation can be admissibly large because the upper bound for β^*, which is a linear function of the population size, is small.

5 Concluding remarks

We examined the diffusion of products and practices with coordination effects through contagion using evolutionary game theory framework. Evolutionary game theory captures many realistic aspects of individual decision making and interactions, most notably, myopia (i.e. the inability to remember the entire history of play in complex social interactions); the tendency to experiment or make mistakes on optimal choices; and the locality of social interactions. We consider two related models that capture these structural and behavioural properties of social interactions: evolution with best response, and evolution with best response and mutations.

Using these two models, we show that a p-dominant strategy of a symmetric coordination game, if it exists, is contagious in networks with the contagion threshold equal or greater than p. We defined a measure of the contagion threshold for finite networks that is easily computable. We then examined the effects of noise on contagion, showing that there exists a threshold level of noise below which contagious choices are long-run equilibrium of an evolutionary process where agents make mistakes.

Our results have broader implications for targeted contagion. First, we provide steps for identifying the smallest set of agents that sufficiently trigger the contagious spread of a p-dominant action. We find that this set can be as small as three agents, and is independent of the population size. Second, our results imply that targeted contagion is economically reasonable even in settings where agents experiment with their choices.

Finally, our results have implications for convergence rates of evolutionary dynamics in networks. One of the common criticisms of the evolutionary game models is that the convergence rates to equilibrium tend to increase with the population size, so that in very large populations, equilibrium is not reached in timescales of economic relevance. In the analysis that we have relegated to Section A, we show that if a contagious strategy exist, then equilibrium is reached fast.

A Expected waiting time

This section focuses on Model 2 to examine how the process of contagion affects the expected waiting time from any state to the state with the highest long-run probability. The problem of slow diffusion does not arise in situations where the level of experimentation is sufficiently high. Kreindler and Young (2013) and Kreindler and Young (2014) show that when β is sufficiently small, convergence is fast. Here, we examine the case where β is very large (i.e. $\beta \to \infty$).

We show that when β is large, the expected waiting time to the monomorphic convention corresponding to the contagious strategy is independent of the population size. The direct implication of this result is that even in large networks and with low levels of experimentation, the diffusion process converges fast to the long-run equilibrium. The expected waiting time is formally defined as follows.

Definition 8. *Let $W \subset \mathbf{X}$ be a subset of the configuration space and $\bar{W}$ its complement. Define $T(W) = \inf\{t \geq 0 \,|\mathbf{x}_t \in W\}$ to be the first time W is reached. The expected waiting time from some configuration $\mathbf{x} \in \bar{W}$ to W is then defined as $\mathbb{E}\left[T(W)\,|\mathbf{x}_0 = \mathbf{x}\right]$.*

Let a^* be a p-dominant strategy of (A, U), and assume that $p \leq \eta(G)$ so that a^* is contagious on G. We aim to show that $(A, U, N, G, P_\varepsilon)$ converges to $\mathbf{a}^*$ *fast* (i.e. the expected waiting time to $\mathbf{a}^*$ is independently of the population size). To do so, we show that if a^* is contagious on G, then there exists a function $F(\beta)$ that is independent of n so that $\mathbb{E}\left[T(\mathbf{a}^*) \,|\, \mathbf{x}_0 = \mathbf{x}\right] \leq F(\beta)$ for all initial configurations $\mathbf{x}_0 \neq \mathbf{a}^*$.

Proposition 2. *Let a^* be a p-dominant strategy of (A, U) and assume that a^* is contagious on a strongly connected network G so that $\mathbf{a}^*$ is the long-run equilibrium of the diffusion process (A, U, N, G, P_β). Then there exists some $b^* \in (0, b_2^*)$ such that*

$$\lim_{\beta \to \infty} \frac{\ln \mathbb{E}[T(\mathbf{a}^*)]}{\beta} = b^* \tag{5}$$

Proof. See Appendix D

$\square$

Proposition 2 shows that the expected waiting time to the equilibrium configuration of the evolutionary process with mutations from any other configuration takes the form

$$\mathbb{E}[T(\mathbf{a}^*)] \leq \exp\left[\beta b^* + g(\beta)\right] = F(\beta) \tag{6}$$

where $g(\beta)$ is a decreasing function of noise parameter β. As the level of experimentation tends to zero (i.e. $\beta \to \infty$), $g(\beta) \to 0$ and the expected waiting time increases at an exponential rate of b^*, which is independent of the population size.

Compared to existing results on convergence rates of evolutionary processes such as Ellison (1993), Young (2011), Kreindler and Young (2013) and Kreindler and Young (2014), the result in Proposition 2 is driven more by contagion and less by noise. Kreindler and Young (2014) also find that learning is fast in networks, but they consider a 2×2 coordination game with random sampling and with deterministic dynamics. Moreover, they define fast learning as the case in which noise is large to the extent that only one unique equilibrium exists. On a contrary, Proposition 2 shows that contagion makes learning fast in stochastic evolutionary game dynamics for $m \times m$ coordination games.

B Proof of Theorem 1

The proof of Theorem 1 follows in two steps. The first step, which is already discussed in detail in Section 3, demonstrates that if $a^* \in A$ is a p-dominant strategy of (A, U) and $p \leq \eta(G)$, then a^* spreads contagiously on G, and that the size of the set of initial adopters is bounded from above by b_2^*. The second step demonstrates that if $a^* \in A$ is a p-dominant strategy and $p \leq \eta(G)$, then the number of mutations needed to leave convention $\mathbf{a}^*$ is greater than $n^{\frac{3}{5}}$.

Strategy a^* spreads contagiously on G: Given the diffusion process (A, U, N, G, P) on a strongly connected network G, and a p-dominant strategy $a^* \in A$, let $p \leq \eta(G)$. Let also $A_1 = A \backslash a^*$. Since a^* is p-dominant, it is a best response when played by at least proportion p of a player's neighbours and the rest play strategies in A_1. This implies that if all players in

a given subgroup $Z \subset N$ play a^*, then a^* is a unique best response to any $i \in N$ for whom $\alpha_i(Z) \geq p$.

Let the evolutionary process start from any state $\mathbf{a} \in \mathbf{L} \backslash a^*$, where $\mathbf{a}$ can consist of only strategies in A_1 or both a^* and strategies in A_1. Pick any $i \in N$ and the respective B_{i_2}, and at period $t = 1$, let all players in B_{i_2} mutate to play a^*. The evolutionary process will evolve through best response from $t = 1$ onward as follows, where we write $\bar{B}_{i_r} = N \backslash B_{i_r}$ for the complement of B_{i_r}, $i \to A_l$ to mean that i plays a strategy in A_l, $Z \to A_l$ to mean that each $j \in Z$ plays a strategy in A_l:

$t = 1$	$B_{i_2} \to a^*$; and $\bar{B}_{i_2} \to A_1 \cup a^*$.
$t = 2$	$B_{i_2} \to a^*$ since each $j \in B_{i_2}$ has $\alpha_j(B_{i_2}) \geq \alpha_i^* \geq p$; $N_{i_3} \to a^*$ since each $j \in N_{i_3}$ has $\alpha_j(B_{i_2}) \geq \alpha_i^* \geq p$; and $\bar{B}_{i_3} \to A_1 \cup a^*$.
$t = 3$	$B_{i_2} \cup N_{i_3} \to a^*$; $N_{i_4} \to a^*$ since each $j \in N_{i_4}$ has $\alpha_j(B_{i_3}) \geq \alpha_i^* \geq p$; and $\bar{B}_{i_4} \to A_1 \cup a^*$.
$t = 4$	$B_{i_3} \cup N_{i_4} \to a^*$; $N_{i_5} \to a^*$ since each $j \in N_{i_5}$ has $\alpha_j(B_{i_4}) \geq \alpha_i^* \geq p$; and $\bar{B}_{i_5} \to A_1 \cup a^*$.
$- - -$	$- - - - -$This iterative process continues until $t = d_i - 1 - - - - --$
$t = d_i - 1$	$N \to a^*$.

We see from the above iterative process that after $t = d_i - 1$ iterations, the evolutionary process converges to convention a^*. Thus, starting from any $\mathbf{a} \in \mathbf{L} \backslash a^*$, there exists a best response sequence $\{\mathbf{x}_t\}_{t=0}^{d_i-1}$ with $\mathbf{x}_0 = \mathbf{a}$ and $x_1^j = a^*$ for all $j \in N(\mathbf{a} \to \mathbf{a}^*) = B_{i_2}$, for any $i \in N$, satisfying $x_{d_i-1}^l = a^*$ for all $l \in N$. Since this holds for B_{i_2} of any $i \in N$, it follows that the size of the smallest set of initial adopters is $b_2^* = \arg \min_{i \in N} b_{i_2}$.

Since b_2^* is independent of the population size, n, we conclude that when $p \leq \eta(G)$, strategy a^* spreads contagiously on G with $n(\mathbf{a} \to \mathbf{a}^*) \leq b_2^*$ for all $\mathbf{a} \in \mathbf{L} \backslash a^*$. Note that the upper bound for $n(\mathbf{a} \to \mathbf{a}^*)$ follows because although b_2^* mutations to a^* guarantee that a^* spreads contagiously, it is possible, in most networks, to find a smaller set than B_2^* that sufficiently triggers the contagious spread of a^*.

Uninvadability of $\mathbf{a}^*$: Recall that convention $\mathbf{a}^*$ is uninvadable if $r(\mathbf{a}^*) > n(\mathbf{a} \to \mathbf{a}^*)$ for all $\mathbf{a} \in \mathbf{L} \backslash a^*$, and that $r(\mathbf{a}^*)$, which is the number of mutations required to leave convention $\mathbf{a}^*$, is a function of n. We first show that if a^* spreads contagiously on a strongly connected network G, then $r(\mathbf{a}^*) \geq n^{\frac{3}{5}}$.

Let $R(\mathbf{a}^*) \subset N$ be the smallest set of players that should mutate to strategies in A_1 for the evolutionary process to leave the basin of attraction of $\mathbf{a}^*$, where $r(\mathbf{a}^*)$ is the cardinality of $R(\mathbf{a}^*)$. We see from the preceding analysis that if there exists a player i in G for whom all players in B_{i_2} play a^*, then a^* spreads contagiously, and hence, the evolutionary process converges to $\mathbf{a}^*$ regardless of the strategy configuration of other players not in B_{i_2}. Thus, to leave $\mathbf{a}^*$, no such player must exist, and that each player must be at most two steps away from $R(\mathbf{a}^*)$; that is, for each i, a *sufficiently large proportion* of players in B_{i_2} are in $R(\mathbf{a}^*)$. Since a^*

is p-dominant, where $p \leq \frac{1}{2}$, strategies in A_1 are best responses only when more than proportion $(1-p) > \frac{1}{2}$ of neighbours play strategies in A_1. For strategies in A_1 to become best responses to players within B_{i_2}, more than $\frac{1}{2}$ of the interactions of players in B_{i_2} must be in $R(\mathbf{a}^*)$; and $R(\mathbf{a}^*)$ must be chosen to satisfy this condition for each $i \in N$.

The identification of the smallest $R(\mathbf{a}^*)$ is then equivalent to the problem of identifying the smallest 2-*monopolies* in graph theory, defined as follows. A player i in network G is said to be 2-*controlled* by the set $Z \subset N$ of players if at least half of the players in B_{i_2} are in Z. The set Z is called a 2-monopoly if it 2-controls every player in the network. Bermond et al (1996, Proposition 4) show that the minimum size of a 2-monopoly on any undirected and strongly connected network of size n is $n^{\frac{3}{5}}$ (see Peleg (2002) for a review of the literature on monopolies and local majorities in networks), and hence, $r(\mathbf{a}^*) > n^{\frac{3}{5}}$. Note that the inequality follows because we require $R(\mathbf{a}^*)$ to contain *more than half* of interactions of B_{i_2} for all $i \in N$ while the definition of a 2-monopoly requires $R(\mathbf{a}^*)$ to contain *at least half* of the interactions of B_{i_2}.

Since $r(\mathbf{a}^*) > n^{\frac{3}{5}}$ is an increasing function of n, it follows that convention $\mathbf{a}^*$ is uninvadable whenever $n^{\frac{3}{5}} \geq n(\mathbf{a} \to \mathbf{a}_l)$ for all $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$. And since $n(\mathbf{a} \to \mathbf{a}_l) \leq b_2^*$ for all $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$, convention $\mathbf{a}^*$ is uninvadable, and hence, also contagious, whenever $b_2^* \leq n^{\frac{3}{5}}$.

C   Proof of Proposition 1

To prove Proposition 1, we first characterize the structure of stationary distributions, and in particular, the ratio, $\frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{y})}$, of stationary distributions of any pair of configurations $\mathbf{x}, \mathbf{y} \in \mathbf{X}$. We use the following results from Freidlin and Wentzell (1984).

Lemma 1. *(Freidlin and Wentzell, 1984, Lemma 3.1). Given a diffusion process P_β, the stationary distribution $\pi_\beta(\mathbf{x})$ of some configuration $\mathbf{x} \in \mathbf{X}$ is given by*

$$\pi_\beta(\mathbf{x}) = \left(\sum_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g) \right) \left(\sum_{\mathbf{y} \in \mathbf{X}} \sum_{g \in \Gamma(\{\mathbf{y}\})} P_\beta(g) \right)^{-1} \tag{7}$$

where the total probability $P_\beta(g)$ associated with each graph g is $P_\beta(g) = \prod_{(\mathbf{z}, \mathbf{y}) \in g} P_\beta(\mathbf{z}, \mathbf{y})$ and $\Gamma(\mathbf{x})$ graphs are defined in Definition 7.

To fully characterise the structure of π_β, we first characterise the structure of transition probabilities $P_\beta(\mathbf{x}, \mathbf{y})$ between pairs of states $\mathbf{x}, \mathbf{y} \in \mathbf{X}$. Recall that $c(\mathbf{x}, \mathbf{y})$ is the number of mutations involved in the direct transition from $\mathbf{x}$ to $\mathbf{y}$. That is, the number of players who choose different strategies in state $\mathbf{y}$ than those chosen in state $\mathbf{x}$, and that their choices are a result of mutations. Employing Assumption 1, the transition probability $P_\varepsilon(\mathbf{x}, \mathbf{y})$ can directly be expressed in terms of $c(\mathbf{x}, \mathbf{y})$ as

$$P_\beta(\mathbf{x}, \mathbf{y}) = \left(\frac{e^{-\beta}}{m} \right)^{c(\mathbf{x}, \mathbf{y})} \left(\frac{m + (1-m)e^{-\beta}}{m} \right)^{n - c(\mathbf{x}, \mathbf{y})} \tag{8}$$

The right hand side of (8) follows because, first, if y^i (i.e. the strategy i plays in configuration $\mathbf{y}$) is not a best response to $\mathbf{x}$ so that $BR_i(y^i; \mathbf{x}) = 0$, then from (3), the probability that i plays

y^i is $\frac{e^{-\beta}}{m}$. Consequently, the probability that $c(\mathbf{x}, \mathbf{y})$ players simultaneously play strategies that are not best responses to $\mathbf{x}$ in configuration $\mathbf{y}$ is $\left(\frac{e^{-\beta}}{m}\right)^{c(\mathbf{x},\mathbf{y})}$.

Second, if y^i is instead a best response to $\mathbf{x}$, so that, from Assumption 1, $BR_i(y^i; \mathbf{x}) = 1$, then i plays y^i with probability

$$\frac{e^{-\beta}}{m} + (1 - e^{-\beta}) = \frac{m + (1-m)e^{-\beta}}{m}$$

Thus, the probability that the remaining $n - c(\mathbf{x}, \mathbf{y})$ players simultaneously play strategies that are best responses to $\mathbf{x}$ is $\left(\frac{m+(1-m)e^{-\beta}}{m}\right)^{n-c(\mathbf{x},\mathbf{y})}$.

Next, we characterise the probabilities of $\Gamma(\mathbf{x})$ graphs. From the definition of $\Gamma(\mathbf{x})$ graphs, every $g \in \Gamma(\mathbf{x})$ spans the entire state space except $\mathbf{x}$, which is the root of the graph. Each $\mathbf{y} \in \mathbf{X} \backslash \mathbf{x}$ has only one arrow emanating from it. Thus, each $g \in \Gamma(\mathbf{x})$ contains a total of $m^n - 1$ directed edges, where m^n is the cardinality of $\mathbf{X}$. Similarly, if we let $d(L_j)$ be the cardinality of $D(L_j)$ (the basin of attraction of L_j), then there are $d(L_j)$ directed edges that originate from states in $D(L_j)$.

Now, let $g(D(L))$ be a subgraph of g consisting of all $d(L_j)$ directed edges that originate from states in $D(L_j)$. Since $D(L_j)$ for all $L_j \subset \mathbf{L}$ are non-overlapping sets, we can rewrite $P_\beta(g)$ as

$$P_\beta(g) = \prod_{L_j \in \mathbf{L}} \prod_{(\mathbf{y},\mathbf{z}) \in g(D(L_j))} P_\beta(\mathbf{y}, \mathbf{z}) \tag{9}$$

We can further subdivide the set of edges of $g(D(L_j))$ into those that involve at least one mutation, denoted by $g(D(L_j); \beta)$, and those whose dynamics are governed solely by best response, denoted by $g(D(\bar{L}_j); \beta)$. That is, for each $\mathbf{x} \in D(L_j)$ and some $\mathbf{y} \neq \mathbf{x}$, a directed edge $(\mathbf{x}, \mathbf{y}) \in g(D(L_j); \beta)$ if $c(\mathbf{x}, \mathbf{y}) > 0$, and $(\mathbf{x}, \mathbf{y}) \in g(D(\bar{L}_j); \beta)$ if $c(\mathbf{x}, \mathbf{y}) = 0$. Using these definitions and notation, $P_\beta(g)$ can be rewritten as

$$P_\beta(g) = \prod_{L_j \in \mathbf{L}} \prod_{(\mathbf{y},\mathbf{z}) \in g(D(L_j))} P_\beta(\mathbf{y}, \mathbf{z})$$

$$= \prod_{L_j \in \mathbf{L}} \prod_{(\mathbf{y},\mathbf{z}) \in g(D(L_j))} \left(\frac{e^{-\beta}}{m}\right)^{c(\mathbf{y},\mathbf{z})} \prod_{i=1}^{n-c(\mathbf{y},\mathbf{z})} \left(\frac{m + (1-m)e^{-\beta}}{m}\right)$$

$$= \prod_{L_j \in \mathbf{L}} \left[\prod_{(\mathbf{y},\mathbf{z}) \in g(D(L_j);\beta)} \left(\frac{e^{-\beta}}{m}\right)^{c(\mathbf{y},\mathbf{z})} \left(\frac{m + (1-m)e^{-\beta}}{m}\right)^{n-c(\mathbf{y},\mathbf{z})} \right.$$

$$\times \left. \prod_{(\mathbf{y},\mathbf{z}) \in g(D(\bar{L}_j);\beta)} \left(\frac{m + (1-m)e^{-\beta}}{m}\right)^n \right] \tag{10}$$

Let $n(g; L_j; \beta)$ be the cardinality of subgraph $g(D(L_j); \beta)$ (i.e. the number of directed edges in subgraph $g(D(L_j); \beta)$) so that the cardinality of $g(D(\bar{L}_j); \beta)$ is $d(L_j) - n(g; L_j; \beta)$. Then (10)

can be rewritten as

$$P_\beta(g) = \prod_{L_j \in \mathbf{L}} \left[\left(\frac{e^{-\beta}}{m} \right)^{\sum_{(\mathbf{y},\mathbf{z}) \in g(D(L_j);\beta)} c(\mathbf{y},\mathbf{z})} \left(\frac{m + (1-m)e^{-\beta}}{m} \right)^{\sum_{(\mathbf{y},\mathbf{z}) \in g(D(L_j);\beta)} (n - c(\mathbf{y},\mathbf{z}))} \right. $$

$$\left. \times \left(\frac{m + (1-m)e^{-\beta}}{m} \right)^{n[d(L_j) - n(g;L_j;\beta)]} \right] \tag{11}$$

$$= \prod_{L_j \in \mathbf{L}} \left(\frac{e^{-\beta}}{m} \right)^{\sum_{(\mathbf{y},\mathbf{z}) \in g(D(L_j);\beta)} c(\mathbf{y},\mathbf{z})} \left(\frac{m + (1-m)e^{-\beta}}{m} \right)^{n[d(L_j) - n(g;L_j;\beta)] + \sum_{(\mathbf{y},\mathbf{z}) \in g(D(L_j);\beta)} (n - c(\mathbf{y},\mathbf{z}))}$$

The summation $c(L_j;g) = \sum_{(\mathbf{y},\mathbf{z}) \in g(D(L_j);\beta)} c(\mathbf{y},\mathbf{z})$ on the exponent of the expressions of $P_\beta(g)$ in (11) is the total cost of leaving the basin of attraction of each L_j under graph $g \in \Gamma(\mathbf{x})$. Using this definition, we can simplify $n[d(L_j) - n(g;L_j;\beta)] + \sum_{(\mathbf{y},\mathbf{z}) \in g(D(L_j);\beta)} (n - c(\mathbf{y},\mathbf{z}))$ as follows:

$$n[d(L_j) - n(g;L_j;\beta)] + \sum_{(\mathbf{y},\mathbf{z}) \in g(D(L_j);\beta)} (n - c(\mathbf{y},\mathbf{z})) = n[d(L_j) - n(g;L_j;\beta)] + n(g;L_j;\beta)n - c(L_j;g)$$

$$= nd(L_j) - c(L_j;g)$$

Equation (11) then simplifies to

$$P_\beta(g) = \prod_{L_j \in \mathbf{L}} \left(\frac{e^{-\beta}}{m} \right)^{c(L_j;g)} \left(\frac{m + (1-m)e^{-\beta}}{m} \right)^{nd(L_j) - c(L_j;g)}$$

$$= \exp\left[-\sum_{L_j \in \mathbf{L}} \left(c(L_j;g)\beta_m + (nd(L_j) - c(L_j;g))\beta'_m \right) \right]$$

$$= \exp\left[-\sum_{L_j \in \mathbf{L}} \left(c(L_j;g)(\beta_m - \beta'_m) + nd(L_j)\beta'_m \right) \right] \tag{12}$$

where $\beta_m = \beta - \ln m^{-1}$ and $\beta'_m = -\ln\left[\frac{m + (1-m)e^{-\beta}}{m} \right]$.

Recall the definition of the long-run equilibrium of $(A, U, N, G, P_\varepsilon)$ as configurations that maximize the stationary distribution. So, to compute the long-run equilibrium, we take ratios of probabilities and identify configurations for which the ratio is less than one. Specifically, configuration $\mathbf{a}^*$ is a long-run equilibrium if $\frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{a}^*)} \leq 1$ for all $\mathbf{x} \neq \mathbf{a}^*$. From $\pi_\beta(\mathbf{x})$ in (7), the expression for the ratio of stationary distribution, $\frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{w})}$ of any pair of configurations $\mathbf{x}, \mathbf{w} \in \mathbf{X}$ is given by

$$\frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{w})} = \frac{\sum_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g)}{\sum_{g \in \Gamma(\{\mathbf{w}\})} P_\beta(g)} \tag{13}$$

where the quantity $\left(\sum_{\mathbf{y} \in \mathbf{X}} \sum_{g \in \Gamma(\{\mathbf{y}\})} P_\beta(g) \right)^{-1}$ cancels out since it is identical for all configurations.

Let $\gamma(\mathbf{x}) = \#\Gamma(\mathbf{x})$ be the cardinality of $\Gamma(\mathbf{x})$. Since for any $\mathbf{x} \in \mathbf{X}$ every $g \in \Gamma(\mathbf{x})$ spans the entire state space except for one configuration (i.e. the number of edges in any $g \in \Gamma(\mathbf{x})$ is equal

to $m^n - 1$), the cardinality of $\Gamma(\mathbf{x})$ is the same for any pair of configurations $\mathbf{x}, \mathbf{y} \in \mathbf{X}$; that is, $\gamma(\mathbf{x}) = \gamma(\mathbf{y}) = \gamma$. Using this notation, the following bounds hold:

$$\sum_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g) \leq \sum_{g \in \Gamma(\{\mathbf{x}\})} \max_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g) = \gamma \max_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g)$$

$$\sum_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g) \geq \max_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g)$$

Thus, the ratio $\frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{w})}$ is bounded from below and above by

$$\frac{\max_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g)}{\gamma \max_{g \in \Gamma(\{\mathbf{w}\})} P_\beta(g)} \leq \frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{w})} \leq \frac{\gamma \max_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g)}{\max_{g \in \Gamma(\{\mathbf{w}\})} P_\beta(g)} \tag{14}$$

Substituting for $P_\beta(g)$ from (12) into the ratios of stationary distributions in (14) yields the following expression.

$$\frac{\max_{g \in \Gamma(\{\mathbf{x}\})} P_\beta(g)}{\max_{g \in \Gamma(\{\mathbf{w}\})} P_\beta(g)} = \frac{\max_{g \in \Gamma(\{\mathbf{x}\})} \exp\left[-\sum_{L_j \in \mathbf{L}} \left(c(L_j; g)(\beta_m - \beta'_m) + nd(L_j)\beta'_m \right) \right]}{\max_{g \in \Gamma(\{\mathbf{w}\})} \exp\left[-\sum_{L_j \in \mathbf{L}} \left(c(L_j; g)(\beta_m - \beta'_m) + nd(L_j)\beta'_m \right) \right]}$$

$$= \exp\left[\min_{g \in \Gamma(\{\mathbf{w}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g)(\beta_m - \beta'_m) - \min_{g \in \Gamma(\{\mathbf{x}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g)(\beta_m - \beta'_m) \right]$$

$$= \exp\left[-(\beta_m - \beta'_m) \left(\min_{g \in \Gamma(\{\mathbf{x}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g) - \min_{g \in \Gamma(\{\mathbf{w}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g) \right) \right] \tag{15}$$

whereby, due to the negative in the exponents, the expressions for $P_\beta(g)$ are maximized when the costs of exiting the basins of attractions of absorbing sets are minimized. Substituting (15) into (14) yields the following lower and upper bounds for $\frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{w})}$

$$\frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{w})} = \exp\left[-(\beta_m - \beta'_m) \left(\min_{g \in \Gamma(\{\mathbf{x}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g) - \min_{g \in \Gamma(\{\mathbf{w}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g) \right) \pm \ln \gamma \right] \tag{16}$$

Now, notice that for all $\beta \geq 0$, $(\beta_m - \beta'_m) \geq 0$: that is,

$$(\beta_m - \beta'_m) = \beta - \ln m^{-1} + \ln\left[\frac{m + (1-m)e^{-\beta}}{m} \right] = \beta - \ln m^{-1} + \ln\left[m + (1-m)e^{-\beta} \right] + \ln m^{-1}$$

$$= \beta + \ln\left[m + (1-m)e^{-\beta} \right] \tag{17}$$

Since $m + (1-m)e^{-\beta} \geq 1$ for all $\beta \geq 0$, it follows from the right hand side of (17) that $(\beta_m - \beta'_m) \geq 0$ for all $\beta \geq 0$. Let $\Phi(\mathbf{x}, \mathbf{w})$ denote the difference between costs of graphs (i.e. the

expression in the exponent of (16)). That is,

$$\Phi(\mathbf{x}, \mathbf{w}) = \min_{g \in \Gamma(\{\mathbf{x}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g) - \min_{g \in \Gamma(\{\mathbf{w}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g)$$

It then follows from (16) that $\frac{\pi_\beta(\mathbf{x})}{\pi_\beta(\mathbf{w})} \leq 1$ if $\Phi(\mathbf{x}, \mathbf{w}) \geq 0$ and that

$$-(\beta_m - \beta'_m)\Phi(\mathbf{x}, \mathbf{w}) + \ln \gamma \leq 0 \tag{18}$$

Substituting for the values of β_m and β'_m in (18) yields

$$\beta \geq \frac{\ln \gamma}{\Phi(\mathbf{x}, \mathbf{w})} - \ln \left[m + (1-m)e^{-\beta} \right] \tag{19}$$

Let β^* be the solution to

$$\beta = \frac{\ln \gamma}{\Phi(\mathbf{x}, \mathbf{w})} - \ln \left[m + (1-m)e^{-\beta} \right] \tag{20}$$

By comparison β^* is smaller than $\frac{\ln \gamma}{\Phi(\mathbf{x},\mathbf{w})}$. To see why, notice that the function $\ln \left[m + (1-m)e^{-\beta} \right]$ increases from 0 to the upper bound $\ln m$ as $\beta \to \infty$, which implies that $\frac{\ln \gamma}{\Phi(\mathbf{x},\mathbf{w})} - \ln \left[m + (1-m)e^{-\beta} \right]$ decreases from $\frac{\ln \gamma}{\Phi(\mathbf{x},\mathbf{w})}$ to $\frac{\ln \gamma}{\Phi(\mathbf{x},\mathbf{w})} - \ln m$. This implies that the equilibrium value of β in (20), that is, the value of β at which the 45-degree line depicting $\beta = \beta$ meets $\frac{\ln \gamma}{\Phi(\mathbf{x},\mathbf{w})} - \ln \left[m + (1-m)e^{-\beta} \right]$, is less than $\frac{\ln \gamma}{\Phi(\mathbf{x},\mathbf{w})}$, and hence,

$$\beta^* \leq \frac{\ln \gamma}{\Phi(\mathbf{x}, \mathbf{w})} \tag{21}$$

It then follows that configuration $\mathbf{a}^*$ is the long-run equilibrium for all $\beta \geq \beta^*$ if it hase the least cost $\mathbf{a}^*$-tree, that is, $\Phi(\mathbf{x}, \mathbf{a}^*) \geq 0$ for all $\mathbf{x} \neq \mathbf{a}^*$.

We now show that if a p-dominant strategy, a^*, is contagious in a given network, then convention $\mathbf{a}^*$ has the minimum cost $\Gamma(\mathbf{a}^*)$ graph. Note that configurations with the minimum cost graph belong to some absorbing state, and hence, it suffices to focus on examining the costs of graphs for configurations within $\mathbf{L}$. Let $\phi(\mathbf{a}, g)$ denote the total cost of some $g \in \Gamma(\mathbf{a})$ for any $\mathbf{a} \in \mathbf{L}$; that is,

$$\phi(\mathbf{a}; g) = \min_{g \in \Gamma(\{\mathbf{a}\})} \sum_{L_j \in \mathbf{L}} c(L_j; g) \tag{22}$$

We can see from (22) that $\phi(\mathbf{a}; g)$ is identical to the cost derived from a reduced form of $\Gamma(\mathbf{a})$, denoted by $\Gamma'(\mathbf{a})$, on a state space $\mathbf{L} \times \mathbf{L}$ (i.e. where vertices are absorbing sets). For any $g \in \Gamma'(\mathbf{a})$ the cost of an arrow originating from some $L_j \in \mathbf{L}$ is $c(L_j; g)$.

We deduce from the proof of Theorem 1 in Appendix (B) that for convention $\mathbf{a}^*$, the minimum cost graph in $\Gamma'(\mathbf{a}^*)$ involves direct transitions from every $\mathbf{a} \in L_j \subset \mathbf{L} \backslash \mathbf{a}^*$ to $\mathbf{a}^*$, and each has a cost bounded from above by $c(\mathbf{a}; g) \leq b_2^*$. Note that if $c(\mathbf{a}; g) \leq b_2^*$ for all $\mathbf{a} \in L_j$, then $c(L_j; g) \leq b_2^*$. Thus, if we let $\zeta^*(\mathbf{L})$ be the number of independent absorbing sets in $\mathbf{L}$ (i.e. all $L_j \subset \mathbf{L}$) with $\mathbf{a}^*$ excluded, then

$$\phi(\mathbf{a}^*) = \min_{g \in \Gamma'(\mathbf{a}^*)} \phi(\mathbf{a}^*; g) \leq b_2^* \zeta^*(\mathbf{L}) \tag{23}$$

Now, consider any other $\mathbf{a} \in L \subset \mathbf{L} \backslash \mathbf{a}^*$. The minimum cost graph in $\Gamma'(\mathbf{a})$ will consist of an arrow $\mathbf{a}^* \to \mathbf{a}'$ starting from $\mathbf{a}^*$ to some $\mathbf{a}' \neq \mathbf{a}^*$. The minimum cost $g \in \Gamma'(\mathbf{a})$ can thus be constructed from the minimum cost graph in $\Gamma'(\mathbf{a}^*)$ by deleting $\mathbf{a} \to \mathbf{a}^*$ and replacing it with $\mathbf{a}^* \to \mathbf{a}'$, so that

$$\phi(\mathbf{a}) = \phi(\mathbf{a}^*) - \min_{g \in \Gamma'(\mathbf{a}^*)} c(\mathbf{a}; g) + \min_{g \in \Gamma'(\mathbf{a})} c(\mathbf{a}^*; g) \tag{24}$$

And hence,

$$\phi(\mathbf{a}) - \phi(\mathbf{a}^*) = \min_{g \in \Gamma'(\mathbf{a})} c(\mathbf{a}^*; g) - \min_{g \in \Gamma'(\mathbf{a}^*)} c(\mathbf{a}; g) > \left(n^{\frac{3}{5}} - b_2^* \right) \tag{25}$$

where the second inequality is because the minimum cost of leaving $\mathbf{a}^*$ is greater than $n^{\frac{3}{5}}$ mutations (see Appendix (B)) and that $c(\mathbf{a}; g) \leq b_2^*$ for all $\mathbf{a} \in L_j \subset \mathbf{L} \backslash \mathbf{a}^*$.

Since $n^{\frac{3}{5}} \geq b_2^*$ is a condition for strategy a^* to be contagious, it follows that $\phi(\mathbf{a}) - \phi(\mathbf{a}^*) > 0$ for all $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$, and hence, $\mathbf{a}^*$ has the minimum cost $\mathbf{a}^*$-tree. The inequality, $\phi(\mathbf{a}) - \phi(\mathbf{a}^*) > n^{\frac{3}{5}} - b_2^*$, also implies that $\Phi(\mathbf{a}, \mathbf{a}^*) > n^{\frac{3}{5}} - b_2^*$ for all $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$. Thus,

$$\frac{\ln \gamma}{\Phi(\mathbf{a}, \mathbf{a}^*)} < \frac{\ln \gamma}{n^{\frac{3}{5}} - b_2^*} \quad \text{for all } \mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$$

It then follows that there exists some $\beta^* \in \left(0, \ln \gamma / (n^{\frac{3}{5}} - b_2^*) \right]$ such that for all $\beta \geq \beta^*$, convention $\mathbf{a}^*$ has the least cost tree, and hence, the long-run equilibrium.

D Proof of Proposition 2

The following definitions are used in the next steps of the proof. Given transition probabilities $P_\beta(\mathbf{x}, \mathbf{y})$ in (8), the cost function $c(\mathbf{x}, \mathbf{y})$ can be rewritten as follows

$$\lim_{\beta \to \infty} \frac{- \ln P_\beta(\mathbf{x}, \mathbf{y})}{\beta} = c(\mathbf{x}, \mathbf{y}) \tag{26}$$

The following definition is for $\Gamma_{\mathbf{x}, \mathbf{y}}(W)$ graphs, which are a special form of $\Gamma(W)$ graphs defined in Definition 7.

Definition 9. *For any $\mathbf{x} \in \bar{W}$ and $\mathbf{y} \in W$ where $\mathbf{x} \neq \mathbf{y}$, $\Gamma_{\mathbf{x}, \mathbf{y}}(W)$ is a set of all $\Gamma(W)$-graphs which link $\mathbf{x}$ to $\mathbf{y}$. For any two configurations $\mathbf{x}, \mathbf{y} \in \bar{W}$, $\Gamma_{\mathbf{x}, \mathbf{y}}(W \cup \{\mathbf{y}\})$ is the set of $\Gamma(W)$-graphs in which $\mathbf{x}$ is joined to some point $\mathbf{y}$ possibly itself and not to W, and that all other points of $\bar{W}$ are joined to either the same point or to W.*

Consider a configuration space $\mathbf{X} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}, \mathbf{e}, \mathbf{f}, \mathbf{g}, \mathbf{h}\}$ with examples of $\mathbf{g}$-trees depicted in Figure 4. Let $W = \{\mathbf{d}, \mathbf{e}, \mathbf{f}, \mathbf{g}, \mathbf{h}\}$, with examples of $\Gamma(W)$ graphs depicted in Figure 5. Then examples of $\Gamma_{\mathbf{a}, \mathbf{c}}\left(W \cup \{\mathbf{c}\} \right)$ graphs based on $\Gamma(W)$ graphs in Figure 5 are: $\{\mathbf{a} \to \mathbf{c}, \mathbf{b} \to W\}$ for the graph on the left, $\{\mathbf{c} \to W, \mathbf{b} \to W\}$ for the middle graph, and $\{\mathbf{a} \to \mathbf{b}, \mathbf{c} \to W\}$ for the graph on the right.

Recall the definition of $\phi(W; g)$ for some $W \subset \mathbf{X}$ and $g \in \Gamma(W)$ from Section C as $\phi(W; g) = \sum_{(\mathbf{x}, \mathbf{y}) \in g} c(\mathbf{x}, \mathbf{y})$. The following result is derived in Catoni (1999, Proposition 4.2).
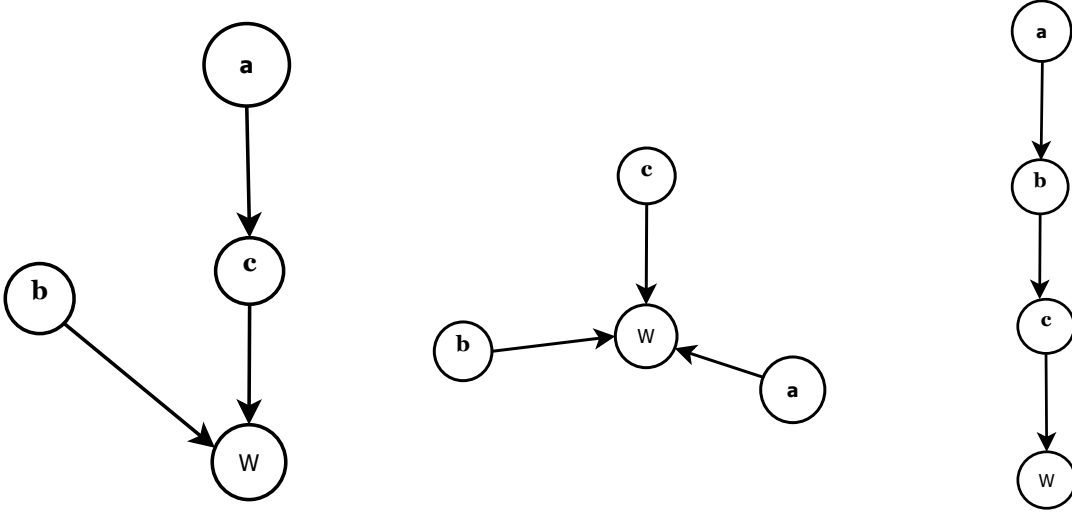
Figure 5: Examples of $\Gamma(W)$ graphs, where $\bar{W} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}\}$.

Lemma 2. *For any $W \subset \mathbf{X}$, $W \neq \emptyset$ and $\bar{W} = \mathbf{X} \backslash W$, for any $\mathbf{x}, \mathbf{y} \in \bar{W}$*

$$\lim_{\beta \to \infty} \frac{\ln \mathbb{E}[T(W) | \mathbf{x}_0 = \mathbf{x}]}{\beta} = \min_{g \in \Gamma(W)} \phi(W; g) - \min_{\mathbf{y} \in \bar{W}} \min_{g \in \Gamma_{\mathbf{x}, \mathbf{y}}(W \cup \{\mathbf{y}\})} \phi(W; g) \tag{27}$$

We are interested in the expected waiting time for convention $\mathbf{a}^*$ associated with the contagious strategy a^*. Thus, we can substitute $W = \{\mathbf{a}^*\}$ into (27). That is,

$$\lim_{\beta \to \infty} \frac{\ln \mathbb{E}[T(\mathbf{a}^*) | \mathbf{x}_0 = \mathbf{x}]}{\beta} = \min_{g \in \Gamma(\mathbf{a}^*)} \phi(\mathbf{a}^*; g) - \min_{\mathbf{y} \in \mathbf{L} \backslash \mathbf{a}^*} \min_{g \in \Gamma_{\mathbf{x}, \mathbf{y}}(\mathbf{a}^* \cup \{\mathbf{y}\})} \phi(\mathbf{a}^*; g) \tag{28}$$

Recall from the analysis in Sections B and C for the minimum cost graph $g \in \Gamma(\mathbf{a}^*)$, the costs are bounded from above as $c(\mathbf{a}; g) \leq b_2^*$ for all $\mathbf{a} \in \mathbf{L} \backslash \mathbf{a}^*$, and from (23), $\phi(\mathbf{a}^*) = \min_{g \in \Gamma(\mathbf{a}^*)} \phi(\mathbf{a}^*; g) \leq b_2^* \zeta^*(\mathbf{L})$. Focusing on reduced form graphs, $\Gamma'(\mathbf{a}^*)$, it follows from the definition of $\Gamma'_{\mathbf{a}, \mathbf{a}'}(\mathbf{a}^* \cup \{\mathbf{a}'\})$ graphs that $\min_{\mathbf{a}' \in \mathbf{L} \backslash \mathbf{a}^*} \min_{g \in \Gamma'_{\mathbf{a}, \mathbf{a}'}(\mathbf{a}^* \cup \{\mathbf{a}'\})} \phi(\mathbf{a}^*; g) \leq b_2^* (\zeta^*(\mathbf{L}) - 1)$. Thus, there exists some $b^* \in (0, b_2^*)$ such that

$$\lim_{\beta \to \infty} \frac{\ln \mathbb{E}[T(\mathbf{a}^*) | \mathbf{x}_0 = \mathbf{a}]}{\beta} = \min_{g \in \Gamma(\mathbf{a}^*)} \phi(\mathbf{a}^*; g) - \min_{\mathbf{a}' \in \mathbf{L} \backslash \mathbf{a}^*} \min_{g \in \Gamma'_{\mathbf{a}, \mathbf{a}'}(\mathbf{a}^* \cup \{\mathbf{a}'\})} \phi(\mathbf{a}^*; g) = b^* \tag{29}$$

Acknowledgements

References

Alós-Ferrer C, Weidenholzer S (2007) Partial bandwagon effects and local interactions. Games and Economic Behavior 61(2):179–197

Alós-Ferrer C, Weidenholzer S (2008) Contagion and efficiency. Journal of Economic Theory 143(1):251–274

Anderson SP, Goeree JK, Holt CA (2001) Minimum-effort coordination games: Stochastic potential and logit equilibrium. Games and Economic Behavior 34(2):177–199

Bass FM (1969) A new product growth for model consumer durables. Management science 15(5):215–227

Beaman L, BenYishay A, Magruder J, Mobarak AM (2018) Can network theory-based targeting increase technology adoption? Tech. rep., National Bureau of Economic Research

Bermond JC, Bond J, Peleg D, Perennes S (1996) Tight bounds on the size of 2-monopolies. In: SIROCCO, pp 170–179

Berninghaus SK, Schwalbe U (1996) Conventions, local interaction, and automata networks. Journal of Evolutionary Economics 6(3):297–312

Blume LE (1995) The statistical mechanics of best-response strategy revision. Games and Economic Behavior 11(2):111–145

Campbell A (2013) Word-of-mouth communication and percolation in social networks. American Economic Review 103(6):2466–98

Catoni O (1999) Simulated annealing algorithms and markov chains with rare transitions. In: Azéma J, Émery M, Ledoux M, Yor M (eds) Séminaire de probabilités XXXIII, Springer, Berlin Heidelberg, pp 69–119

Egidi M, Marris RL, Viale R, et al (1992) Economics, bounded rationality and the cognitive revolution. Edward Elgar Publishing

Ellison G (1993) Learning, local interaction, and coordination. Econometrica 61(5):1047–1071

Ellison G (2000) Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. Review of Economic Studies 67(1):17–45

Ellison G (2006) Bounded rationality in industrial organization. Econometric Society Monographs 42:142

Foster D, Young P (1990) Stochastic evolutionary game dynamics. Theoret Population Biol

Freidlin M, Wentzell AD (1984) Random perturbations of dynamical systems. New York: Springer Verlag

Galeotti A, Goyal S (2009) Influencing the influencers: a theory of strategic diffusion. The RAND Journal of Economics 40(3):509–532

Galeotti A, Goyal S, Jackson MO, Vega-Redondo F, Yariv L (2010) Network games. Review of Economic Studies 77(1):218–244

Goyal S, Heidari H, Kearns M (2014) Competitive contagion in networks. Games and Economic Behavior

Jackson MO, Yariv L (2007) Diffusion of behavior and equilibrium properties in network games. American Economic Review 97(2):92–98

Kandori M, Mailath GJ, Rob R (1993) Learning, mutation, and long run equilibria in games. Econometrica 61(1):29–56

Kreindler GE, Young HP (2013) Fast convergence in evolutionary equilibrium selection. Games and Economic Behavior 80:39 – 67

Kreindler GE, Young HP (2014) Rapid innovation diffusion in social networks. Proceedings of the National Academy of Sciences 111(Supplement 3):10881–10888

Lee IH, Valentinyi A (2000) Noisy contagion without mutation. Review of Economic Studies 67(1):47–56

Lee IH, Szeidl A, Valentinyi A (2003) Contagion and State Dependent Mutations. The BE Journal of Theoretical Economics 3(1):1–29

López-Pintado D, et al (2008) Diffusion in complex social networks. Games and Economic Behavior 62(2):573–590

McKelvey RD, Palfrey TR (1995) Quantal response equilibria for normal form games. Games and economic behavior 10(1):6–38

Morris S (2000) Contagion. Review of Economic Studies 67(1):57–78

Morris S, Rob R, Shin HS (1995) p-Dominance and Belief Potential. Econometrica 63(1):145–57

Opolot D (2018) On the relationship between p-dominance and stochastic stability in network games. Available at SSRN 3234959

Opolot D (2020) Contagion, network cohesion and long run stability in evolutionary games. Available at SSRN 3678191

Oyama D, Takahashi S (2015) Contagion and uninvadability in local interaction games: The bilingual game and general supermodular games. Journal of Economic Theory 157:100–127

Peleg D (2002) Local majorities, coalitions and monopolies in graphs: a review. Theoretical Computer Science 282(2):231–257

Peski M (2010) Generalized risk-dominance and asymmetric dynamics. Journal of Economic Theory 145(1):216–248

Rogers EM (2003) Diffusion of innovations. 5th Edition. Free Press

Sundararajan A (2007) Local network effects and complex network structure. The BE Journal of Theoretical Economics 7(1)

Tsakas N (2017) Diffusion by imitation: The importance of targeting agents. Journal of Economic Behavior & Organization 139:118–151

Young HP (1993) The evolution of conventions. Econometrica 61(1):57–84

Young HP (1998) Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton University Press

Young HP (2009) Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning. American economic review 99(5):1899–1924

Young PH (2011) The dynamics of social innovation. Proceedings of the National Academy of Sciences 108(4):21285–21291