

Enforcing Social Norms: Trust-building and community enforcement*

Joyee Deb[†]

New York University

Julio González-Díaz[‡]

University of Santiago de Compostela

April 3, 2014

Abstract

We study impersonal exchange, and ask how agents can behave honestly in anonymous transactions, without contracts. We analyze repeated anonymous random matching games, where agents observe only their own transactions. Little is known about cooperation in this setting beyond prisoner's dilemma. We show that cooperation can be sustained quite generally, using community enforcement and “trust-building”. The latter refers to an initial phase of the game in which one community builds trust by not deviating despite a short-run incentive to cheat; the other community reciprocates trust by not punishing deviations during this phase. Trust-building is followed by cooperative play, sustained through community enforcement.

* **Acknowledgements.** We are grateful to Johannes Hörner for many insightful discussions and comments. We thank Mehmet Ekmekci, Jérôme Renault, Larry Samuelson, Satoru Takahashi and three anonymous referees for their comments. We also thank seminar participants at Boston University, GAMES 2008, Erice 2010, Northwestern University, New York University, Penn State University, and the Repeated Games workshop at SUNY StonyBrook for many suggestions. The second author gratefully acknowledges the support of the Sixth Framework Programme of the European Commission through a Marie Curie fellowship, and the Ministerio de Ciencia e Innovación through a Ramón y Cajal fellowship and through projects ECO2008-03484-C02-02 and MTM2011-27731-C03. Support from Xunta de Galicia through project INCITE09-207-064-PR is also acknowledged. A version of this paper has been circulated under the title “Community Enforcement Beyond the Prisoner's Dilemma.”

[†]Email: joyee.deb@nyu.edu

[‡]Email: julio.gonzalez@usc.es

1 Introduction

In many economic settings, impersonal exchange in society is facilitated by legal institutions that can enforce contractual obligations. However, there is also a wide range of economic examples where impersonal exchange exists in the absence of legal contractual enforcement –either because existing legal institutions are weak, or even by design. For instance, Greif (1994, 2006) documents impersonal exchange or trade between different communities in medieval Europe before the existence of centralized legal institutions. In large part, the diamond industry, even today, opts out of the legal system (Bernstein, 1992). Diamond dealers have developed norms or rules that enforce honest trading, and the community punishes opportunistic behavior by appropriate sanctions. Millions of transactions are completed on eBay and other internet sites, where buyers and sellers trade essentially anonymously (Resnick et al., 2000). These settings raise the important question of how agents achieve cooperative outcomes, and act in good faith in transactions with strangers, in the absence of formal contracts. This is the central question of our paper.

We model impersonal exchange as an infinitely repeated random matching game, in which players from two different communities are randomly and anonymously matched to each other to play a two-player game. Each player observes only his own transactions: He does not receive any information about the identity of his opponent or about how play proceeds in other transactions. In this setting of “minimal information-transmission,” we ask what payoffs can be achieved in equilibrium. In particular, can agents be prevented from behaving opportunistically? Can they achieve cooperative outcomes that could not be achieved in a one-shot interaction? In practice, cooperation or honest behavior within and between communities is often sustained by *community enforcement*; a mechanism in which a player’s dishonest behavior against one partner is punished by other members of society. Community enforcement arises as a norm in many settings. In this paper, we examine whether such social norms (cooperative behavior enforced by social sanctions) can be sustained in equilibrium with rational agents.

Two early papers by Kandori (1992) and Ellison (1994) showed that community enforcement can sustain cooperative behavior for the special case in which the agents are engaged in a prisoner’s dilemma game (PD). However, it is an open question whether cooperation can be sustained in other strategic situations beyond the PD. In the aforementioned papers, if a player ever faces a defection, he punishes all future rivals by switching to defection forever (Nash reversion). By starting to defect, he spreads the information that

someone has defected. The defection action spreads throughout the population: More and more people get *infected*, and cooperation eventually breaks down completely. The credible threat of such a breakdown of cooperation can deter players from defecting in the first place. However, these arguments rely critically on properties of the PD. In particular, since the Nash equilibrium of the PD is in strictly dominant strategies, the punishment action is dominant and so gives a current gain even if it lowers continuation payoffs. In an arbitrary game, on facing a deviation for the first time, players may not have the incentive to punish, because punishing can both lower future continuation payoffs and entail a short-term loss in that period. How then can cooperation be sustained in general? We establish that it is, indeed, possible to sustain a wide range of payoffs in equilibrium in a large class of games beyond the PD, provided that all players are sufficiently patient and that the population size is not too small. We argue that communities of around twenty people should suffice for our results in most games.

In particular, we show that for any stage-game with a strict Nash equilibrium, the ideas of community enforcement coupled with “trust-building” can be used to sustain cooperation. In equilibrium, play proceeds in two phases. There is an initial phase of what we call “trust-building,” followed by a cooperative phase that lasts forever, as long as nobody deviates. If anyone observes a deviation in the cooperative phase, he triggers Nash reversion (or community enforcement). In the initial trust-building phase, players of one community build trust by not deviating from the equilibrium action even though they have a short-run incentive to do so, and players in the other community reciprocate the trust by not starting any punishments during this phase even if they observe a deviation. This initial trust-building phase turns out to be crucial to sustaining cooperation in the long-run. The ideas that profitable long-term relationships must start by building trust, and that a breach of trust is subsequently punished harshly by the community are very intuitive, and our paper can be viewed as providing a foundation for how such norms of behavior are sustained.¹

To the best of our knowledge, this is the first paper to sustain cooperation in a random matching game beyond the PD without adding any extra informational assumptions. Some papers that go beyond the PD introduce verifiable information about past play to sus-

¹There is a literature on building trust in repeated interactions (see for instance, Ghosh and Ray (1996) and Watson (2002)). The role of trust in this paper has a different flavor. These papers focus on the “gradual” building of trust, where the stakes in a relationship grow over time. Our equilibrium does not feature this gradualism. Rather, we have an initial phase in which players behave cooperatively even though they have an incentive to deviate, and this initial phase is exactly what helps sustain cooperation in the long-run. In this sense, trust is built so that trigger strategies can be effective in sustaining cooperation.

tain cooperation.² More recently, Deb (2012) obtains a general folk theorem by allowing transmission of unverifiable information (cheap talk).³

An important feature of our equilibrium is that the strategies are simple and plausible. Unlike recent work on games with imperfect private monitoring (Ely and Välimäki, 2002; Piccione, 2002; Ely, Hörner and Olszewski, 2005; Hörner and Olszewski, 2006) and, specifically, in repeated random matching games (Takahashi, 2010; Deb, 2012), we do not rely on belief-free ideas. The strategies give the players strict incentives on and off the equilibrium path. Further, unlike existing literature, our strategies are robust to changes in the discount factor.⁴

The reader may wonder how our equilibrium strategies compare with social norms that arise in reality. Our strategies use community enforcement, which has the defining feature that opportunistic behavior by any agent is punished by (possibly all) members of the community, not necessarily by the victim alone. This is indeed a common feature in practice. For instance, in the diamond industry dishonest traders are punished through sanctions by all other community members. Greif (1994, 2006) documents the existence of a community responsibility system in which a player who acted dishonestly was punished by all members of the rival community, who in addition also punished the rest of the deviator's community. In practice, however, societies also try to disseminate information and publicize the identity of deviators for better enforcement. For instance, Bernstein (1992) reports that the diamond dealer's club "facilitates the transmission of information about dealers' reputations, and serves both a reputation-signaling and a reputation-monitoring role." Similarly, eBay and Amazon have developed reputation rating systems to enable transmission

²For instance, Kandori (1992) assumes the existence of a mechanism that assigns labels to players based on their history of play. Players who have deviated or have seen a deviation can be distinguished from those who have not, by their labels. This naturally enables transmission of information, and cooperation can be sustained in a specific class of games. For related approaches, see Dal Bó (2007), Hasker (2007), Okuno-Fujiwara and Postlewaite (1995), and Takahashi (2010).

³There is a recent literature on repeated games and community enforcement on networks. It may be worthwhile to investigate if the ideas of our construction apply in these settings. Yet, there are important modeling differences, which change the incentive issues substantively. On a network players are not anonymous. Moreover, this literature mostly restricts attention to the PD, and partners do not change over time (see, for instance, Ali and Miller (2012) and Lippert and Spagnolo (2011)). More recently, Nava and Piccione (Forthcoming) consider cooperation on a network that can change over time and show that cooperation can be sustained through some notion of reciprocity, which is not a possibility in our anonymous setting.

⁴In a very recent sequence of papers, Sugaya (2013*a,b,c*) establishes general folk theorems under imperfect private monitoring, for two-player PD, general two-player games and N -player games. However, our anonymous random matching setting, even if viewed as an N -player game of imperfect private monitoring, does not fall within the scope of these papers, since it violates the full-support monitoring assumption, as well as other identifiability assumptions required in these papers.

of information about past transactions (Resnick et al., 2000). We do not allow dissemination of information within the community, but rather model all transactions as anonymous (even after deviations). We do this for two reasons. First, the anonymous setting with minimal information transmission offers a useful benchmark: Establishing the possibility of cooperation in this environment tells us that cooperation should certainly be sustainable when more verifiable information is available, and punishments can be personalized by using information about past deviations. Second, the anonymous setting allows us to model applications in which it may be difficult to know the identity of the rival. Consider, for instance, settings in which individual stakes are low but, because the number of agents is large, equilibrium behavior may still have a high social impact: For example, think of the repeated interaction between citizens in a district and the staff that keeps the district streets clean. Actions are taken asynchronously and it is natural to assume that the agents cannot transmit information about each other. Citizens can choose to litter or not. If the streets are too dirty, *i.e.*, if the citizens litter, then a staff effort might not have an incentive to exert effort, since such an effort would not make a difference. In this setting, we can ask if it is possible to sustain an equilibrium in which citizens do not litter, and the staff exerts effort, when the outcome in a static setting might be for the citizens to litter and for the staff to shirk.

It is worth emphasizing that this paper also makes a methodological contribution, since we develop techniques to work explicitly with players' beliefs. We use Markov chains to model the beliefs of the players *off the equilibrium path*. We hope that the methods we use to study the evolution of beliefs will be of independent interest, and can be applied in other contexts such as the study of belief-based equilibria in general repeated games.

It is useful to explain why working with beliefs is fundamental to our approach. Recall that the main challenge to sustaining cooperation through community enforcement is that, when an agent is supposed to punish a deviation, he may find that doing so is costly for both his current and his continuation payoffs. The main feature of our construction with trust-building is that, when a player is required to punish by reverting to the Nash action, his off-path beliefs are such that he thinks that most people are already playing Nash. To see how this works, start by assuming that players entertained the possibility of correlated deviations. Then, we could assume that, upon observing a deviation, a player thinks that all the players in the rival community have simultaneously deviated and that everybody is about to start the punishment. This would make Nash reversion optimal, but it is an arti-

ficial way to guarantee a common flow of information and get coordinated punishments.⁵ Indeed, if we rule out the possibility of coordinated deviations, a player who observes a deviation early in the game will know that few players have been affected so far and that Nash reversion will not be optimal for him. This suggests that, to induce appropriate beliefs, equilibrium strategies must prescribe something different from Nash reversion in the initial periods of the game, which is the reason for the trust-building phase.⁶

The rest of the paper is organized as follows. Section 2 contains the model and the main result. In Section 3, we illustrate the result using a leading example of the product-choice game. This section is useful in understanding the strategies and the intuition behind the result. In Section 4, we discuss the generality of our result and potential extensions. We provide the general equilibrium construction and the formal proof in the Appendix.

2 Cooperation Beyond the PD

2.1 The Setting

There are $2M$ players, divided into two communities with M players each. In each period $t \in \{1, 2, \dots\}$, players are randomly matched into pairs, with each player from Community 1 facing a player from Community 2. The matching is anonymous, independent and uniform over time.⁷ After being matched, each pair of players plays a finite two-player game G . The action set of each player $i \in \{1, 2\}$ is denoted by A_i and an action profile is an element of $A := A_1 \times A_2$. Players observe the actions and payoffs in their match. Then, a new matching occurs in the next period. Throughout, we refer to arbitrary players and players in Community 1 as male players and to players in Community 2 as female players.

Players can observe only the transactions they are personally engaged in, *i.e.*, each player knows the history of action profiles played in each of his stage-games in the past. A player never observes his opponent's identity. Further, he gets no information about how other players have been matched or about the actions chosen by any other pair of

⁵The solution concept used in this paper is sequential equilibrium (Kreps and Wilson, 1982), which rules out such off-path beliefs. In contrast such beliefs would be admissible, for instance, under weak perfect Bayesian equilibrium.

⁶Ideally, we would like strategies such that every player, at each information set, has a unique best reply that is independent of his beliefs (as in Kandori (1992) and Ellison (1994)). We have not been able to construct such strategies.

⁷Although the assumption of uniform matching greatly simplifies the calculations, we expect our results to hold for other similar matching technologies.

players. All players have discount factor $\delta \in (0, 1)$, and their payoffs in the infinitely repeated random matching game are the normalized sum of the discounted payoffs from the stage-games. No public randomization device is assumed.⁸ The solution concept used is sequential equilibrium. Given a game G , the associated repeated random matching game with communities each of size M and discount factor δ is denoted by G_δ^M .⁹

2.2 A negative result

The main difficulty in sustaining cooperation through community enforcement is that players may not have the the incentive to punish deviations. Below, we present a simple example to show that a straightforward adaptation of grim trigger strategies (or the contagion strategies as in Ellison (1994)) cannot be used to support cooperation in general.

		Buyer	
		B_H	B_L
Seller	Q_H	$1, 1$	$-l, 1 - c$
	Q_L	$1 + g, -l$	$0, 0$

Figure 1: The product-choice game.

Consider the simultaneous-move game between a buyer and a seller presented in Figure 1. Suppose that this game is played by a community of M buyers and a community of M sellers in the repeated anonymous random matching setting. In each period, every seller is randomly matched with a buyer. After being matched, each pair plays the product-choice game, where $g > 0$, $c > 0$, and $l > 0$.¹⁰ The seller can exert either high effort (Q_H) or low effort (Q_L) in the production of his output. The buyer, without observing the seller's choice, can buy either a high-priced product (B_H) or a low-priced product (B_L). The buyer prefers the high-priced product if the seller has exerted high effort and prefers the low-priced product if the seller has not. For the seller, exerting low effort is a dominant action. The efficient outcome of this game is (Q_H, B_H) , while the Nash equilibrium is (Q_L, B_L) . We denote a product-choice game by $\Gamma(g, l, c)$. The infinitely repeated random matching

⁸Section 4 has a discussion of what can be gained with a randomization device.

⁹Note that we are assuming that there are two independent communities. Alternatively, we could have taken one population whose members are matched in pairs in every period and, in each match, the roles of players are randomly assigned. We believe that the trust-building ideas that underlie this paper can be adapted to this new setting and refer the reader to Online Appendix (B.7) for a discussion. It is worth noting that the choice to model independent communities is driven by the fact that it seems to be the hardest scenario to achieve cooperation because of the negative result in Section 2.2.

¹⁰See Mailath and Samuelson (2006) for a discussion of this game in the context of repeated games.

game associated with the product-choice game $\Gamma(g, l, c)$, with discount parameter δ and communities of size M , is denoted by $\Gamma_\delta^M(g, l, c)$.

Proposition 1. *Let $\Gamma(g, l, c)$ be a product-choice game with $c \leq 1$. Then, there is $\underline{M} \in \mathbb{N}$ such that, for each $M \geq \underline{M}$, regardless of the discount factor δ , the repeated random matching game $\Gamma_\delta^M(g, l, c)$ has no sequential equilibrium in which (Q_H, B_H) is played in every period on the equilibrium path, and in which players play the Nash action off the equilibrium path.*

Proof. Suppose that there exists an equilibrium in which (Q_H, B_H) is played in every period on the equilibrium path, and in which players play the Nash action off the equilibrium path. Suppose that a seller deviates in period 1. We argue below that, for a buyer who observes this deviation, it is not optimal to switch to the Nash action permanently from period 2. In particular, we show that playing B_H in period 2, followed by switching to B_L from period 3 onwards gives the buyer a higher payoff. The buyer who observes the deviation knows that, in period 2, with probability $\frac{M-1}{M}$ she will face a different seller who will play Q_H . Consider this buyer's short-run and long-run incentives:

Short-run: The buyer's payoff in period 2 from playing B_H is $\frac{1}{M}(-l) + \frac{M-1}{M}$. Her payoff if she switches to B_L is $\frac{M-1}{M}(1-c)$. Hence, if M is large enough, she has no short-run incentive to switch to the Nash action.

Long-run: With probability $\frac{1}{M}$, the buyer will meet the deviant seller (who is already playing Q_L) in period 2. In this case, her action will not affect this seller's future behavior, and so her continuation payoff will be the same regardless of her action.

With probability $\frac{M-1}{M}$, the buyer will meet a different seller. Note that, since $1 - c \geq 0$, a buyer always prefers to face a seller playing Q_H . So, regardless of the buyer's strategy, the larger the number of sellers who have already switched to Q_L , the lower is her continuation payoff. Hence, playing B_L in period 2 will give her a lower continuation payoff than playing B_H , because action B_L will make a new seller switch permanently to Q_L .

Since there is no short-run or long-run incentive to switch to the Nash action in period 2, the buyer will not start punishing. Therefore, playing (Q_H, B_H) in every period on-path, and playing the Nash action off-path does not constitute a sequential equilibrium, regardless of the discount factor. \square

Notice that the product-choice game represents a minimal departure from the PD. If we replace the payoff $1 - c$ with $1 + g$, we get the PD. However, even with this small departure from the PD, cooperation can not be sustained in equilibrium using the standard grim trigger strategies.¹¹ In the next section, we state our main result which is a possibility result about cooperation in a class of games beyond the PD.

2.3 The Main Result

CLASS OF GAMES: We denote by \mathcal{G} the class of finite two-player games with the following two properties:

P1. Strict Nash equilibrium. There exists a strict Nash equilibrium a^* .¹²

P2. One-sided Incentives. There exists a pure action profile (\hat{a}_1, \hat{a}_2) in which one player has a strict incentive to deviate while the other has a strict incentive to stick to the current action.

ACHIEVABLE PAYOFFS: Let G be a game and let $\underline{a} \in A$. Let $A_{\underline{a}} := \{a \in A : a_1 = \underline{a}_1 \iff a_2 = \underline{a}_2\}$. Define $F_{\underline{a}} := \text{conv}\{u(a) : a \in A_{\underline{a}}\} \cap \{v \in \mathbb{R}^2 : v > u(\underline{a})\}$.

Our main result says that given a game G in \mathcal{G} with a strict Nash equilibrium a^* , it is possible for players to achieve any payoff in F_{a^*} in equilibrium in the corresponding infinitely repeated random matching game G_{δ}^M , if players are sufficiently patient and the communities are not too small. \mathcal{G} is a large class of games and, in particular, includes the PD. Before stating the result formally, we discuss the assumptions P1 and P2.

Since we are interested in Nash reversion, the existence of a pure Nash equilibrium is needed.¹³ The extra requirement of strictness eliminates only games that are non-generic. To see why we need P1, note that under imperfect private monitoring, it is typically impossible to get full coordination in the punishments. Therefore, if a player is asked to revert to a Nash equilibrium that is not strict, he may not be willing to do. That is, if he thinks that there is a positive probability that his opponent will not play the Nash action, he may

¹¹It is worth emphasizing that while Proposition 1 establishes that cooperation cannot be sustained via grim trigger strategies, it does not rule out the possibility of sustaining cooperation using other repeated game strategies. Establishing such a negative result for patient players for the full class of repeated game strategies, although desirable, is typically very challenging, and we have not been able to obtain such a result.

¹²A Nash equilibrium is strict if each player has a unique best response to his rival's actions (Harsanyi, 1973). By definition, a strict equilibrium is in pure actions.

¹³In the literature on repeated random matching games, it is common to restrict attention to pure strategies.

not be indifferent between two actions in the support of his Nash action and, indeed, his myopic best reply may even be outside the support.

P2 is a mild condition on the class of games. \mathcal{G} excludes what we call *games with strictly aligned interests*: for two-player games this means that, at each action profile, a player has a strict incentive to deviate if and only if his opponent also has a strict incentive to deviate. The games in \mathcal{G} are generic in the class of games without strictly aligned interests with a pure Nash equilibrium.¹⁴

The set of achievable payoffs includes payoffs arbitrarily close to efficiency for the prisoner’s dilemma. In general, our strategies do not suffice to get a folk theorem for games in \mathcal{G} . However, we conjecture that, by adequately modifying our strategies, it may be possible to support payoffs outside F_{a^*} and obtain a Nash threats folk theorem for the games in \mathcal{G} .¹⁵ We now present the formal statement of the result.

Theorem 1. *Let G be a game in \mathcal{G} with a strict Nash equilibrium a^* . There exists $\underline{M} \in \mathbb{N}$ such that, for each payoff profile $v \in F_{a^*}$, each $\varepsilon > 0$, and each $M \geq \underline{M}$, there exists $\underline{\delta} \in (0, 1)$ such that there is a strategy profile in the repeated random matching game G_δ^M that constitutes a sequential equilibrium for each $\delta \in [\underline{\delta}, 1)$ with payoff within ε of v .*

A noteworthy feature of our equilibrium strategies is that if a strategy profile constitutes an equilibrium for a given discount factor, it does so for any higher discount factor as well; in particular, the equilibrium strategy profile defines what is called a *uniform equilibrium* (Sorin, 1990). This is in contrast with existing literature, where strategies have to be fine-tuned based on the discount factor (e.g. Takahashi (2010) and Deb (2012)).¹⁶

Interestingly, the continuation payoff is within ε of the target equilibrium payoff v , not just ex-ante, but throughout the game on the equilibrium path. In this sense, cooperation is sustained as a durable phenomenon, which contrasts with reputation models where, for every δ , there exists a time after which cooperation collapses (e.g. Cripps, Mailath and Samuelson (2004)).

¹⁴We have not been able to apply our approach to games of strictly aligned interests. We refer the reader to the Online Appendix (B.9) for an example that illustrates the difficulty with achieving cooperation in certain games in this class. However, cooperation is not an issue in commonly studied games in this class, like “battle of the sexes” and the standard version of “chicken,” since in these games, the set of Pareto efficient payoffs is spanned by the set of pure Nash payoffs (so, we can alternate the pure Nash action profiles with the desired frequencies).

¹⁵We refer the reader to the Online Appendix (B.8) for a discussion.

¹⁶Further, in Ellison (1994), the severity of punishments depends on the discount factor which has to be common for all players. We just need all players to be sufficiently patient.

While cooperation with a larger population calls for more patient players (larger δ), a very small population also hinders cooperation. Our construction requires a minimum community size M . We work explicitly with beliefs off-path, and a relatively large M guarantees that the beliefs induce the correct incentives to punish. The lower bound \underline{M} on the community size depends only on the game G ; it is independent of the target payoff v and the precision ε . In sections A.3.3 and A.4 in the Appendix we discuss rough bounds on M that suggest that populations of size 20 would be enough for many games.¹⁷

2.4 Equilibrium Strategies

Let G be a game in \mathcal{G} and let a^* be a strict Nash equilibrium. Let (\hat{a}_1, \hat{a}_2) be a pure action profile in which only one player has an incentive to deviate. Without loss of generality, suppose it is player 1 who wants to deviate from (\hat{a}_1, \hat{a}_2) while player 2 does not and let \bar{a}_1 be the most profitable deviation. Let the target equilibrium payoff be $v \in F_{a^*}$. We maintain the convention that players 1 and 2 of the stage-game belong to communities 1 and 2, respectively. Also, henceforth, when we refer to the Nash action, we mean the strict Nash equilibrium action profile a^* . We present the strategies that sustain v . We divide the game into three phases (Figure 2). Phases I and II are *trust-building* phases and Phase III is the *target payoff* phase.

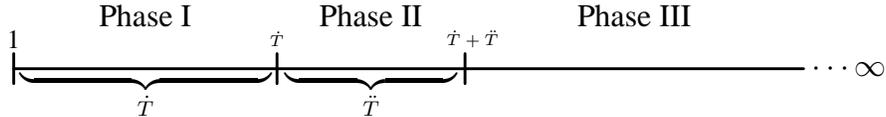


Figure 2: Different phases of the strategy profiles.

Equilibrium play: Phase I: During the first \bar{T} periods, action profile (\hat{a}_1, \hat{a}_2) is played. In every period in this phase, players from Community 1 have a short-run incentive to deviate, but those from Community 2 do not. **Phase II:** During the next \bar{T} periods, players play (a_1^*, a_2) , an action profile where players from Community 1 play their Nash action and players from Community 2 do not play their best response.¹⁸ In every period in this phase, players from Community 2 have a short-run incentive to deviate. **Phase III:** For the rest of the game, the players play a sequence of pure

¹⁷It would be interesting to investigate if our strategies are robust to entry and exit of players. We think that our equilibrium construction may be extended to a setting in which we allow players to enter and exit in each period with small probability, but such an analysis is beyond the scope of this paper.

¹⁸Player 2's action a_2 can be any action other than a_2^* in the stage-game.

action profiles that approximates the target payoff v and such that no player plays according to a^* in period $\hat{T} + \ddot{T} + 1$.

Off-Equilibrium play: A player can be in one of two moods: *uninfected* or *infected*, with the latter mood being irreversible. At the start of the game, all players are uninfected. We classify deviations into two types of actions. Any deviation by a player from Community 2 in Phase I is called a *non-triggering* action. Any other deviation is a *triggering* action. A player who has observed a triggering action is in the infected mood. An uninfected player continues to play as if on-path.¹⁹ An infected player acts as follows.

- A player who gets infected after facing a triggering action switches to the Nash action forever, either from the beginning of Phase II or immediately from the next period, whichever is later.²⁰ In other words, any player in Community 2 who faces a triggering action in Phase I switches to her Nash action forever from the start of Phase II, playing as if on-path in the meantime. A player facing a triggering action at any other stage of the game immediately switches to the Nash action forever.
- A player who gets infected by playing a triggering action himself henceforth best responds, period by period, to the strategies of the other players. In particular, we will show, for large enough \hat{T} , a player from Community 1 who deviates in the first period will continue to (profitably) deviate until the end of Phase I, and then switch to playing the Nash action forever.²¹

Note that a profitable deviation by a player is punished (ultimately) by the whole community of players, with the punishment action spreading like an epidemic. We refer to the spread of punishments as contagion.

The difference between our strategies and standard contagion (Kandori, 1992; Ellison, 1994) is that here, the game starts with two trust-building phases. This can be interpreted as follows. In Phase I, players in Community 1 build credibility by not deviating even

¹⁹In other words, if a player from Community 1 observes a deviation in Phase I, he ignores the deviation, and continues to play as if on-path.

²⁰Moreover, if such an infected player decides to deviate himself in the future, he will subsequently play ignoring his own deviation.

²¹In particular, even if the deviating player observes a history that has probability zero, he will continue to best respond given his beliefs (described in Section A.2).

though they have a short-run incentive to do so. The situation is reversed in Phase II, where players in Community 2 build credibility by not playing a_2^* , even though they have a short-run incentive to do so. A deviation by a player in Community 1 in Phase I is not punished in his trust-building phase, but is punished as soon as the phase is over. Similarly, if a player in Community 2 deviates in her trust-building phase, she effectively faces punishment once the trust-building phase is over. Unlike the results for the PD, where the equilibria are based on trigger strategies, we have “delayed” trigger strategies. In Phase III, deviations immediately trigger Nash reversion.

Next, we discuss the main ideas underlying the proof of Theorem 1 for the particular case of the product-choice game. The general proof can be found in the Appendix.

3 An Illustration: The Product-Choice Game

We show that, for the product-choice game (Figure 1), it is possible to sustain the efficient action profile in equilibrium, using strategies as described above in Section 2.4.

3.1 Equilibrium Strategies in the Product-Choice Game

Equilibrium play: **Phase I:** During the first \dot{T} periods, (Q_H, B_H) is played. **Phase II:** During the next \ddot{T} periods, the players play (Q_L, B_H) . **Phase III:** For the rest of the game, the players play the efficient action profile (Q_H, B_H) .

Off-Equilibrium play: Action B_L in Phase I is a *non-triggering* action. Any other deviation is a *triggering* action. An uninfected player continues to play as if on-path. A player who gets infected after facing a triggering action switches to the Nash action forever either from the beginning of Phase II or immediately from the next period, whichever is later. A player facing a triggering action at any other stage immediately switches to the Nash action forever. A player who gets infected by playing a triggering action best responds to the strategies of the other players.

Clearly, the payoff from the strategy profile above will be arbitrarily close to the efficient payoff $(1, 1)$ for δ large enough. We need to establish that the strategy profile constitutes a sequential equilibrium of the repeated random matching game $\Gamma_\delta^M(g, l, c)$ when \dot{T} , \ddot{T} , and M are appropriately chosen and δ is close enough to 1.

3.2 Optimality of Strategies: Intuition

The incentives on-path are quite straightforward. Any short-run profitable deviation will eventually trigger Nash reversion that will spread and reduce continuation payoffs. Hence, given \dot{T} , \ddot{T} , and M , for sufficiently patient players, the future loss in continuation payoff will outweigh any current gain from deviation. Establishing sequential rationality off-path is the challenge. Below, we consider some key histories and argue why the strategies are optimal after these histories. We start with two observations.

First, a seller who deviates to make a short-term gain in period 1 will find it optimal to revert to the Nash action immediately. This seller knows that, regardless of his choice of actions, from period \dot{T} on, at least one buyer will start playing Nash. Then, from period $\dot{T} + \ddot{T}$ on, more and more players will start playing the Nash action, and contagion will spread exponentially fast. Thus, his continuation payoff after $\dot{T} + \ddot{T}$ will be quite low, regardless of what he does in the remainder of Phase I. Our construction is such that if \dot{T} is large enough, both in absolute terms and relative to \ddot{T} , no matter how patient this seller is, the best thing he can do after deviating in period 1 is to play the Nash action forever.²²

Second, the optimal action of a player after he observes a triggering action depends on the beliefs that he has about how far the contagion has already spread. To see why, think of a buyer who observes a triggering action during, say, Phase III. Is Nash reversion optimal for her? If she believes that few people are infected, then playing the Nash action may not be optimal. In such a case, with high probability she will face a seller playing Q_H and playing the Nash action will entail a loss in that period. Moreover, she is likely to infect her opponent, hastening the contagion and lowering her own continuation payoff. The situation is different if she believes that almost everybody is infected (so, playing Nash). Then, there is a short-run gain by playing the Nash action in this period. Moreover, the effect on the contagion process and the continuation payoff will be negligible. Since the optimal action for a player after observing a triggering action depends on the beliefs he has about “how widespread the contagion is,” we need to define a system of beliefs and check if Nash reversion is optimal after getting infected, given these beliefs.

We make the following assumptions on beliefs. If an uninfected player observes a non-

²²For the deviant seller’s incentives, the choice of large \dot{T} with $\dot{T} \gg \ddot{T}$ is important for two reasons: First, if $\dot{T} \gg \ddot{T}$, then a seller who deviates in period 1 will find it optimal to keep deviating and making short-run profits in Phase I, without caring about potential losses in Phase II. Second, as long as $\dot{T} \gg \ddot{T}$, the deviant seller will believe that the contagion is spread widely enough in Phase I that he will be willing to play Nash throughout Phase III, regardless of the history he observes.

triggering action, then he just thinks that the opponent made a mistake and that no one is infected. If a player observes a triggering action, he thinks that some seller deviated in period 1 and that contagion has been spreading since then. Such beliefs are natural given the interpretation of our strategies: deviations become less likely as more trust is built. Here, we formally study the most extreme such beliefs, which leads to cleaner proofs. In the Online Appendix (B.1) we prove that they are consistent in the sense required by sequential equilibrium.

These beliefs, along with the fact that a deviant seller will play the Nash action forever, imply that, for sufficiently large \dot{T} , any player who observes a triggering action thinks that, almost everybody must be infected by the end of Phase I. This makes Nash reversion optimal after the end of Phase I. To gain some insight, consider the following histories.

- **Suppose that I am a buyer who gets infected in Phase I.** I think that a seller deviated in the first period and that he will continue infecting buyers throughout Phase I. Unless M is too small, in each of the remaining periods of Phase I, the probability of meeting the same seller again is low; so, I prefer to play B_H during Phase I (since all other sellers are playing Q_H). Yet, if \dot{T} is large enough, once Phase I is over I will think that, with high probability, every buyer is now infected. Nash reversion thereafter is optimal.

It may happen that, after I get infected, I observe Q_L in most (possibly all) periods of Phase I. Then, I think that I met the deviant seller repeatedly, and so not all buyers are infected. However, it turns out that if \ddot{T} is large enough, I will still revert to Nash play. Since I expect my continuation payoff to drop after $\dot{T} + \ddot{T}$ anyway, for \ddot{T} large enough, I prefer to play the myopic best reply in Phase II, to make some short-term gains (similar to the argument for a seller's best reply if he deviates in period 1).²³

- **Suppose that I am a seller who faces B_L in Phase I.** (Non-triggering actions) Since such actions are never profitable (on-path or off-path), after observing such an action, I will think it was a mistake and that no one is infected. Then, it is optimal to ignore it. The deviating player knows this, and so it is also optimal for him to ignore it.

- **Suppose that I am a player who gets infected in Phase II, or shortly after period $\dot{T} + \ddot{T}$.** I know that contagion has been spreading since the first period. But the fact that I was uninfected so far may indicate that, possibly, not so many people are infected. We show that if \dot{T} is large enough and $\dot{T} \gg \ddot{T}$, I will still think that, with high probability, I

²³The reason why we need Phase II is to account for the histories in which the deviant seller meets the same buyer in most of the periods of Phase I. Since the probability of these histories is extremely low, they barely affect the incentives of a seller, and so having \ddot{T} periods of Phase II is probably not necessary. Yet, introducing Phase II is convenient since it allows us to completely characterize off-path behavior of the buyers.

was just lucky not to have been infected so far, but that everybody is infected now. This makes Nash reversion optimal.

• **Suppose that I get infected late in the game, at period $\bar{t} \gg \dot{T} + \ddot{T}$.** If $\bar{t} \gg \dot{T} + \ddot{T}$, it is no longer possible to rely on how large \dot{T} is to characterize my beliefs. However, for this and other related histories late in the game, it turns out that, if M is reasonably large, I will still believe that “enough” people are infected and already playing the Nash action, so that playing the Nash action is also optimal for me.

3.3 Choice of Parameters for Equilibrium Construction

It is important to clarify how the different parameters, M , \ddot{T} , \dot{T} , and $\underline{\delta}$, are chosen to construct the equilibrium. First, given a game, we find M so that i) a buyer who is infected in Phase I does not revert to the Nash action before Phase II, and ii) players who are infected very late in the game believe that enough people are infected for Nash reversion to be optimal. Then, we choose \ddot{T} so that, in Phase II, any infected buyer will find it optimal to revert to the Nash action (even if she observed Q_L in all periods of Phase I). Then, we pick \dot{T} , with $\dot{T} \gg \ddot{T}$, so that i) players infected in Phase II or early in Phase III believe that almost everybody is infected, and ii) a seller who deviates in period 1 plays the Nash action ever after. Finally, we pick $\underline{\delta}$ large enough so that players do not deviate on path.²⁴

The role of the discount factor δ requires further explanation. A high δ deters players from deviating from cooperation, but also makes players want to slow down the contagion. Then, why are very patient players willing to spread the contagion after getting infected? The essence of the argument is that, for a fixed M , once a player is infected he knows that contagion will start spreading, and expected future payoffs will converge to 0 exponentially fast. Indeed, we can derive an upper bound on the future gains any player can make by slowing down the contagion once it has started, even if he were perfectly patient ($\delta = 1$). Once we know that the gains from slowing down the contagion are bounded in this way, it is easy to show that even very patient players will be willing to play the Nash action.

²⁴Note that $\underline{\delta}$ must also be large enough so that the payoff achieved in equilibrium is close enough to $(1, 1)$.

4 Discussion and Extensions

4.1 Introduction of Noise

An equilibrium is said to be globally stable equilibrium if, after any finite history, play finally reverts to cooperative play (Kandori, 1992). The notion is appealing because it implies that a single mistake does not entail permanent reversion to punishments. The equilibrium here fails to satisfy global stability, though this can be obtained by introducing a public randomization device.²⁵

A related question is to see if cooperation can be sustained in a model with some noise. Since players have strict incentives, our equilibria are robust to the introduction of some noise in the payoffs. Suppose, however, that we are in a setting where players are constrained to make mistakes with probability at least $\varepsilon > 0$ (small) at every possible history. Our equilibrium construction is not robust to this modification. The incentive compatibility of our strategies relies on the fact that players believe that early deviations are more likely. This ensures that whenever players are required to punish, they think that the contagion has spread enough for punishing to be optimal. If players make mistakes with positive and equal probability in all periods, this property is lost. However, we conjecture that, depending on the game, suitable modifications of our strategies may be robust to the introduction of mistakes. To see why, think of a player who observes a deviation for the first time late in Phase III. In a setting in which mistakes occur with probability ε at every period, this player may not believe that the contagion is widespread, but will still think that it has spread to a certain degree. Whether this is enough to provide him with the incentive to punish, depends on the specific payoffs of the stage-game.²⁶

²⁵This is similar to Ellison (1994). The randomization device could be used to allow for the possibility of restarting the game in any period, with a low but positive probability.

²⁶It is worth noting that if we consider the same strategies we used to prove Theorem 1, there is one particularly problematic situation. Consider the product-choice game in a setting with noise. If a buyer makes a mistake late in Phase II, no matter what she does after that, she will start Phase III knowing that not many people are already infected. Hence, if she is very patient, it may be optimal for her to continue play as if on-path and slow down the contagion. Suppose that a seller observes a triggering action in the last period of Phase II. Since sellers do not spread the contagion in Phase II, this seller will think that it is very likely that his opponent was uninfected and has just made a mistake, and so will not punish. In this case, neither player reverts to Nash punishments. This implies that a buyer may profitably deviate in the last period of Phase II, since her deviation would go unpunished.

4.2 Uncertainty about Timing

In the equilibrium strategies in this paper, players condition behavior on calendar time. On-path, players in Community 1 switch their action in a coordinated way at the end of Phases I and II. Off-path, players coordinate the start of the punishment phase. The calendar time and timing of phases (\dot{T} and \ddot{T}) are commonly known and used to coordinate behavior. Arguably, in modeling large communities, the need to switch behavior with precise coordination is an unappealing feature. It may be interesting to investigate if cooperation can be sustained if players are not sure about the precise time to switch phases.

A complete analysis of this issue is beyond the scope of this paper. We conjecture that a modification of our strategies would be robust to the introduction of *small* uncertainty about timing. The reader may refer to the Online Appendix (B.6), where we consider an altered environment in which players are slightly uncertain about the timing of the different phases. We conjecture equilibrium strategies in this setting, and provide the intuition behind them.

4.3 Alternative Systems of Beliefs

We assume that a player who observes a triggering action believes that some player from Community 1 deviated in the first period of the game. This ensures that an infected player thinks that the contagion has been spreading long enough that, after Phase I, almost everybody is infected. It is easy to see that alternative (less extreme) assumptions on beliefs would still have delivered this property. We work with this case mainly for tractability. Also, since our equilibrium is based on communities building trust in the initial phases of the game, it is plausible that players regard deviations to be more likely earlier rather than later.

Further, the assumption we make is a limiting one in the sense that it yields the weakest bound on M . With other assumptions, for a given game $G \in \mathcal{G}$ and given \dot{T} and \ddot{T} , the threshold population size \underline{M} required to sustain cooperation would be weakly greater than the threshold we obtain. Why is this so? On observing a triggering action, my belief about the number of infected people is determined by my belief about when the first deviation took place and the subsequent contagion process. Formally, on getting infected at period t , my belief x^t can be expressed as $x^t = \sum_{\tau=1}^t \mu(\tau) y^t(\tau)$, where $\mu(\tau)$ is the probability I assign to the first deviation having occurred at period τ , and $y^t(\tau)$ is my belief about the number of people infected if I know that the first deviation took place at period τ . Since contagion is not reversible, every elapsed period of contagion results in a weakly greater

number of infected people. Thus, my belief if I think the first infection occurred at $t = 1$ first-order stochastically dominates my belief if I think the first infection happened later, at any $t > 1$, *i.e.*, for each τ and each $l \in \{1, \dots, M\}$, $\sum_{i=l}^M y_i^t(1) \geq \sum_{i=l}^M y_i^t(\tau)$. Now consider any belief \hat{x}^t that I might have had with alternative assumptions on when I think the first deviation occurred. This belief will be some convex combination of the beliefs $y^t(\tau)$, for $\tau = 1, \dots, t$. Since we know that $y^t(1)$ first-order stochastically dominates $y^t(\tau)$ for all $\tau > 1$, it follows that $y^t(1)$ will also first-order stochastically dominate \hat{x}^t . Therefore, our assumption is the one for which players will think that the contagion is most widespread at any given time and so makes the off-path incentives easier to satisfy.

References

- Ali, Nageeb, and David Miller.** 2012. “Efficiency in repeated games with local interaction and uncertain local monitoring.”
- Bernstein, Lisa.** 1992. “Opting Out of the Legal System: Extralegal Contractual Relations in the Diamond Industry.” *Journal of Legal Studies*, 115.
- Cripps, Martin W., George J. Mailath, and Larry Samuelson.** 2004. “Imperfect Monitoring and Impermanent Reputations.” *Econometrica*, 72(2): 407–432.
- Dal Bó, P.** 2007. “Social norms, cooperation and inequality.” *Economic Theory*, 30(1): 89–105.
- Deb, J.** 2012. “Cooperation and Community Responsibility: A Folk Theorem for Random Matching Games with Names.” mimeo.
- Ellison, Glenn.** 1994. “Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching.” *Review of Economic Studies*, 61(3): 567–88.
- Ely, J. C., and J. Välimäki.** 2002. “A Robust Folk Theorem for the Prisoner’s Dilemma.” *Journal of Economic Theory*, 102(1): 84–105.
- Ely, Jeffrey C., Johannes Hörner, and Wojciech Olszewski.** 2005. “Belief-Free Equilibria in Repeated Games.” *Econometrica*, 73(2): 377–415.
- Ghosh, Parikshit, and Debraj Ray.** 1996. “Cooperation in Community Interaction without Information Flows.” *The Review of Economic Studies*, 63: 491–519.

- Greif, Avner.** 1994. "Cultural Beliefs and the Organization of Society: A Historical and Theoretical Reflection on Collectivist and Individualist Societies." *Journal of Political Economy*, 102(5): 912–950.
- Greif, Avner.** 2006. *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press.
- Harsanyi, J.** 1973. "Oddness of the number of equilibrium points: A new proof." *International Journal of Game Theory*, 2: 235–250.
- Hasker, K.** 2007. "Social norms and choice: a weak folk theorem for repeated matching games." *International Journal of Game Theory*, 36(1): 137–146.
- Hörner, Johannes, and Wojciech Olszewski.** 2006. "The folk theorem for games with private almost-perfect monitoring." *Econometrica*, 74(6): 1499–1544.
- Kandori, Michihiro.** 1992. "Social Norms and Community Enforcement." *Review of Economic Studies*, 59(1): 63–80.
- Kreps, David M., and Robert Wilson.** 1982. "Sequential Equilibria." *Econometrica*, 50: 863–894.
- Lippert, Steffen, and Giancarlo Spagnolo.** 2011. "Networks of Relations and Word-of-Mouth Communication." *Games and Economic Behavior*, 72(1): 202–217.
- Mailath, George J., and Larry Samuelson.** 2006. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press.
- Nava, Francesco, and Michele Piccione.** Forthcoming. "Efficiency in repeated games with local interaction and uncertain local monitoring." *Theoretical Economics*.
- Okuno-Fujiwara, M., and A. Postlewaite.** 1995. "Social Norms and Random Matching Games." *Games and Economic Behavior*, 9: 79–109.
- Piccione, M.** 2002. "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring." *Journal of Economic Theory*, 102(1): 70–83.
- Resnick, Paul, Richard Zeckhauser, Eric Freidman, and Ko Kuwabara.** 2000. "Reputation Systems: Facilitating Trust in Internet Interactions." *Communications of the ACM*, 43(12): 45–48.

- Seneta, Eugene.** 2006. *Non-negative Matrices and Markov Chains. Springer Series in Statistics*, Springer.
- Sorin, S.** 1990. “Supergames.” In *Game theory and applications.*, ed. T. Ichiishi, A. Neyman and Y. Taumann, 46–63. Academic Press.
- Sugaya, Takuo.** 2013a. “Folk Theorem in Repeated Games with Private Monitoring.”
- Sugaya, Takuo.** 2013b. “Folk Theorem in Repeated Games with Private Monitoring with More than Two Players.”
- Sugaya, Takuo.** 2013c. “Folk Theorem in Two-Player Repeated Games with Private Monitoring.”
- Takahashi, Satoru.** 2010. “Community enforcement when players observe partners’ past play.” *Journal of Economic Theory*, 145: 42–62.
- Watson, Joel.** 2002. “Starting Small and Commitment.” *Games and Economic Behavior*, 38: 176–199.

A Proof of the Main Result

We show below that, under the assumptions of Theorem 1, the strategy profile described in Section 2.4 constitutes a sequential equilibrium of the repeated random matching game G_{δ}^M when M is large enough, \hat{T} and \check{T} are appropriately chosen, and δ is close enough to 1.

If we compare the general game to the product-choice game, a player in Community 1 in a general game in \mathcal{G} is like the seller who is the only player with an incentive to deviate in Phase I, and a player in Community 2 is like the buyer who has no incentive to deviate in Phase I but has an incentive to deviate in Phase II. For the sake of clarity, we often refer to players from Community 1 as sellers and players from Community 2 as buyers.

A.1 Incentives on-path

The incentives on-path are straightforward, and so we omit the formal proof. The logic is that if players are sufficiently patient, on-path deviations can be deterred by the threat of eventual Nash reversion. It is easy to see that, if \hat{T} is large enough, the most “profitable”

on-path deviation is that of a player in Community 1 (a seller) in period 1. Given M , \hat{T} and \ddot{T} , the discount factor $\underline{\delta}$ can be chosen close enough to 1 to deter such deviations.

A.2 System of beliefs

In this section we describe the assumptions we make for the formation of off-path beliefs. As usual, they are then updated using Bayes rule whenever possible.

i) **Assumption 1:** If a player observes a triggering action, then this player believes that some player in Community 1 (a seller) deviated in the first period of the game. Subsequently, his beliefs are described as follows:

- **Assumption 1.1:** He believes that after the deviation by a seller in the first period, play has proceeded as prescribed by the strategies, provided his observed history can be explained in this way.
- **Assumption 1.2:** If his observed history cannot be explained with play having proceeded as prescribed by the strategies after a seller's deviation in period 1, then we call this an *erroneous history*.²⁷ In this case, the player believes that after the seller deviated in the first period of the game, one or more of his opponents in his matches made a mistake. Indeed, this player will think that there have been as many mistakes by his past rivals as needed to explain the history at hand. Erroneous histories include the following:
 - A player observes an action in Phases II or III that is neither part of the strict Nash profile a^* nor the action that is supposed to be played in equilibrium.
 - A player who, after being certain that all the players in the other community are infected, faces an opponent who does not play the Nash action. This can only happen in Phase III. For instance, this happens if an infected player infects $M - 1$ opponents by playing a triggering action against them.

ii) **Assumption 2:** If a player observes a non-triggering action, then this player believes that his opponent made a mistake.

iii) **Assumption 3:** If a player plays a triggering action and then observes a history that has probability zero according to his beliefs, we also say that he is at an erroneous

²⁷For the formation of a player's beliefs after erroneous histories, we assume that "mistakes" by infected players are infinitely more likely than "mistakes" by uninfected players.

history and will think that there have been as many mistakes by his past rivals as needed to explain the history at hand.

In the Online Appendix (B.1) we formally prove consistency of these beliefs.

A.2.1 Modeling Beliefs with Contagion Matrices

So far, we have not formally described the structure of a player's beliefs. The payoff-relevant feature of a player's beliefs is the number of people he believes to be currently infected. Accordingly, we let a vector $x^t \in \mathbb{R}^M$ denote the beliefs of player i about the number of infected people in the other community at the end of period t , where x_k^t denotes the probability he assigns to exactly k people being infected in the other community. To illustrate, when player i observes the first triggering action, Assumption 1 implies that $x^1 = (1, 0, \dots, 0)$. In some abuse of notation, when it is known that a player assigns 0 probability to more than k opponents being infected, we work with $x^t \in \mathbb{R}^k$. We say a belief $x^t \in \mathbb{R}^k$ *first-order stochastically dominates* a belief \hat{x}^t if x^t assigns higher probability to more people being infected; *i.e.*, for each $l \in \{1, \dots, k\}$, $\sum_{i=l}^k x_i^t \geq \sum_{i=l}^k \hat{x}_i^t$. Let \mathcal{I}^t denote the random variable representing the number of infected people in the other community at the end of period t . Let k^t denote the event “ k people in the other community are infected by the end of period t ”, *i.e.*, k^t and $\mathcal{I}^t = k$ denote the same event.

As we will see below, the beliefs of players after different histories evolve according to simple Markov processes, and so can be studied using an appropriate transition matrix and an initial belief. We define below a useful class of matrices: *contagion matrices*. The element c_{ij} of a *contagion matrix* C denotes the probability that the state “ i rivals infected” transitions to the state “ j rivals infected.” Formally, if we let \mathcal{M}_k denote the set of $k \times k$ matrices with real entries, we say that a matrix $C \in \mathcal{M}_k$ is a contagion matrix if it has the following properties:

- i) All the entries of C belong to $[0, 1]$ (represent probabilities).
- ii) C is upper triangular (being infected is irreversible).
- iii) All diagonal entries are strictly positive (with some probability, infected people meet other infected people and contagion does not spread in the current period).
- iv) For each $i > 1$, $c_{i-1,i}$ is strictly positive (With some probability, exactly one new person gets infected in the current period, unless everybody is already infected.).

A useful technical property is that, since contagion matrices are upper triangular, their eigenvalues correspond to the diagonal entries. Given $y \in \mathbb{R}^k$, let $\|y\| := \sum_{i \in \{1, \dots, k\}} y_i$. We will often be interested in the limit behavior of $\frac{yC^t}{\|yC^t\|}$, where C is a contagion matrix and y is a probability vector. Given a matrix C , let $C_{\setminus l}$ denote the matrix derived by removing the last l rows and columns from C . Similarly, $C_{\setminus k}$ is the matrix derived by removing the first k rows and columns and $C_{\setminus k, \setminus l}$ by doing both operations simultaneously. Clearly, if we perform any of these operations on a contagion matrix, we get a new contagion matrix.

A.3 Incentives off-path

To prove sequential rationality, we need to examine players' incentives after possible off-path histories, given the beliefs. This is the heart of the proof and the exposition proceeds as follows. We classify possible off-path histories of a player i based on when player i observed off-path behavior for the first time.

- Given the beliefs described above, it is important to first characterize the best response of a seller who deviates in the first period of the game.
- Next, we consider histories where player i observes a triggering action (gets infected) for the first time in the target payoff phase (Phase III).
- We then consider histories where player i observes a triggering action for the first time during one of the two trust-building phases.
- We discuss non-triggering actions.
- Finally, we discuss two particular types of off-path histories that might be of special interest, since the arguments used here can be applied to other similar histories.

We need some extra notation to characterize the possible histories that can arise in the game. Denote a t -period private history for a player i by h^t . At any time period, actions are classified as two types, depending on the information they convey.

Good behavior (g): We say that a player observes good behavior, denoted by g , if he observes the action an uninfected player would choose. There is one exception here: if the on-path and the off-path actions coincide with the Nash action, then observing the Nash action is good behavior only if the player is uninfected.²⁸

²⁸If a player is infected and observes a Nash action in one such period, this does not convey information about how widespread the contagion is. In particular, it does not point towards fewer people being infected.

Bad behavior (b): We say that a player observes bad behavior, denoted by b , if he does not observe good behavior.²⁹

In a nutshell, *good behavior* points in the direction of fewer people being infected and *bad behavior* does not. If player i observes a t -period history h^t followed by three periods of good behavior and then one period of bad behavior, we represent this by $h^t g g g b$. For any player i , let g_t denote the event “player i observed g in period t ,” and let U^t denote the event “player i is uninfected at the end of period t .” In an abuse of notation, a player’s history is written omitting his own actions. For most of the arguments below, we discuss beliefs from the point of view of a fixed player i , and so often refer to player i in the first person.

A.3.1 Computing off-path beliefs

Since we work with off-path beliefs of players, it is useful to clarify at the outset our approach to computing beliefs. Suppose that I get infected at some period \bar{t} in Phase III and that I face an uninfected player in period $\bar{t} + 1$. I will think that a seller deviated in period 1 and that in period $\bar{T} + \bar{\bar{T}} + 1$ all infected buyers and sellers played the Nash action (which is triggering in this period). Therefore, period $\bar{T} + \bar{\bar{T}} + 2$ must always start with the same number of players infected in both communities. So it suffices to compute beliefs about the number of people infected in the rival community. These beliefs are represented by $x^{\bar{t}+1} \in \mathbb{R}^M$, where $x_k^{\bar{t}+1}$ is the probability of exactly k people being infected after period $\bar{t} + 1$, and must be computed using Bayes rule and conditioning on my private history. What information do I have after history $h^{\bar{t}+1}$? I know that a seller deviated at period 1, so $x^1 = (1, 0, \dots, 0)$. I also know that, after any period $t < \bar{t}$, I was not infected (U^t). Moreover, since I got infected at period \bar{t} , at least one player in the rival community got infected in the same period. Finally, since I faced an uninfected player at $\bar{t} + 1$, at most $M - 2$ people were infected after any period $t < \bar{t}$ (i.e., $\mathcal{I}^t \leq M - 2$).

To compute $x^{\bar{t}+1}$, we compute a series of intermediate beliefs x^t , for $t < \bar{t} + 1$. We compute x^2 from x^1 by conditioning on U^2 and $\mathcal{I}^2 \leq M - 2$; then we compute x^3 from x^2 and so on. Note that, to compute x^2 , we do not use the information that “I did not get infected at any period $2 < t < \bar{t}$.” So, at each $t < \bar{t}$, x^t represents my beliefs when I condition on the fact that the contagion started at period 1 and that no matching that leads to

²⁹By Assumption 1.2, if a player observes off-path behavior that does not correspond with the action an infected player would play, he thinks that this corresponds to a “mistake” by an infected player, pointing in the direction of more people being infected. Thus, this is classified as bad behavior.

more than $M - 2$ people being infected could have been realized.³⁰ Put differently, at each period, I compute my beliefs by eliminating (assigning zero probability to) the matchings I know could not have taken place. At a given period $\tau < \bar{t}$, the information that “I did not get infected at any period t , with $\tau < t < \bar{t}$ ” is not used. This extra information is added period by period, *i.e.*, only at period t we add the information coming from the fact that “I was not infected at period t .” In the Online Appendix (B.3), we show that this method of computing $x^{\bar{t}+1}$ generates the required beliefs, *i.e.*, beliefs at period $\bar{t} + 1$ conditioning on the entire history. Now we are equipped to check the optimality of the equilibrium strategies.

A.3.2 A seller deviates at beginning of the game

The strategies specify that a player who gets infected by deviating to a triggering action henceforth plays his best response to the other players’ strategies. As we argued for the product-choice game, it is easy to show that for a given M , if \dot{T} is large enough (and large relative to \ddot{T}), a seller who deviates from equilibrium action \hat{a}_1 in period 1 will find it optimal to continue deviating and playing his most profitable deviation, \bar{a}_1 , until the end of Phase I to maximize short-run gains, and then will switch to playing the Nash a_1^* forever. Note that, here, it is important that \dot{T} be chosen large enough relative to \ddot{T} . First, it ensures that if the seller keeps deviating in Phase I, his short-run gains in Phase I are larger than the potential losses in Phase II. Second, even if the deviant seller faces many occurrences of a non-Nash action in Phase II, he will assign high probability to the event that all but one buyer got infected in Phase I, and that he has been repeatedly meeting the only uninfected buyer in Phase II. Under such beliefs, playing the Nash action in Phase III is still optimal for him.

A.3.3 A player gets infected in Phase III

Our analysis of incentives in Phase III is structured as follows. First we check the off-path incentives of a player infected at the start of Phase III. Next, we analyze incentives of a player after infection very late in Phase III. Finally, we use a monotonicity argument on the beliefs to check the incentives after infection in intermediate periods in Phase III.

Case 1: Infection at the start of Phase III. Let $h^{\dot{T}+\ddot{T}+1}$ denote a history in which I am

³⁰The updating after period \bar{t} is different, since I know that I was infected at \bar{t} and that no more than $M - 1$ people could possibly be infected in the other community at the end of period \bar{t} .

a player who gets infected in period $\dot{T} + \ddot{T} + 1$. The equilibrium strategies prescribe that I switch to the Nash action forever. For this to be optimal, I must believe that enough players in the other community are already infected. My beliefs depend on what I know about how contagion spreads in Phase I after a seller deviates in period 1. In Phase I, only one deviant seller is infecting buyers (recall that buyers cannot spread contagion in Phase I). The contagion, then, is a Markov process with state space $\{1, \dots, M\}$, where the state represents the number of infected buyers. The transition matrix corresponds with the contagion matrix $S_M \in \mathcal{M}_M$, where a state k transits to $k + 1$ if the deviant seller meets an uninfected buyer, which has probability $\frac{M-k}{M}$. With the remaining probability, *i.e.*, $\frac{k}{M}$, state k remains at state k . When no confusion arises, we omit subscript M in matrix S_M . Let s_{kl} be the probability that state k transitions to state l . Then,

$$S_M = \begin{pmatrix} \frac{1}{M} & \frac{M-1}{M} & 0 & 0 & \dots & 0 \\ 0 & \frac{2}{M} & \frac{M-2}{M} & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & & \vdots \\ 0 & 0 & 0 & \frac{M-2}{M} & \frac{2}{M} & 0 \\ 0 & 0 & 0 & 0 & \frac{M-1}{M} & \frac{1}{M} \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

The transition matrix represents how the contagion is expected to spread. To compute my beliefs after being infected, I must also condition on the information from my own history. Consider any period $t < \dot{T}$. After observing history $h^{\dot{T}+\ddot{T}+1} = g \dots gb$, I know that at the end of period $t + 1$, at most $M - 1$ buyers were infected and I was not infected. Therefore, to compute x^{t+1} , my intermediate beliefs about the number of buyers who were infected at the end of period $t + 1$ (*i.e.*, about \mathcal{I}^{t+1}), I need to condition on the following:

- i) My beliefs about \mathcal{I}^t : x^t .
- ii) I was uninfected at the end of $t + 1$: the event U^{t+1} (this is irrelevant if I am a seller, since sellers cannot get infected in Phase I).
- iii) At most $M - 1$ buyers were infected by the end of period $t + 1$: $\mathcal{I}^{t+1} \leq M - 1$ (otherwise I could not have been uninfected at the start of Phase III).

Therefore, given $l < M$, if I am a buyer, the probability that exactly l buyers are infected after period $t + 1$, conditional on the above information, is given by:

$$\begin{aligned} \mathbb{P}(l^{t+1} | x^t \cap U^{t+1} \cap \mathcal{I}^{t+1} \leq M - 1) &= \frac{\mathbb{P}(l^{t+1} \cap U^{t+1} \cap \mathcal{I}^{t+1} \leq M - 1 | x^t)}{\mathbb{P}(U^{t+1} \cap \mathcal{I}^{t+1} \leq M - 1 | x^t)} \\ &= \frac{x_{l-1}^t s_{l-1,l} \frac{M-l}{M-l+1} + x_l^t s_{l,l}}{\sum_{k=1}^{M-1} (x_{k-1}^t s_{k-1,k} \frac{M-k}{M-k+1} + x_k^t s_{k,k})}. \end{aligned}$$

The expression for a seller would be analogous, but without the $\frac{M-l}{M-l+1}$ factors. Notice that we can express the transition from x^t to x^{t+1} using what we call the *conditional transition matrix*. Let $C \in \mathcal{M}_M$ be defined, for each pair $k, l \in \{1, \dots, M - 1\}$, by $c_{kl} := s_{kl} \frac{M-l}{M-k}$; by $c_{MM} = 1$, and with all remaining entries being 0.

Since we already know that $x_M^t = x_M^{t+1} = 0$, we can just work in \mathbb{R}^{M-1} . Recall that C_{\lfloor} and S_{\lfloor} denote the matrices obtained from C and S by removing the last row and the last column of each. So, the truncated matrix of conditional transitional probabilities C_{\lfloor} is as follows:

$$C_{\lfloor} = \begin{pmatrix} \frac{1}{M} & \frac{M-1}{M} \frac{M-2}{M-1} & 0 & 0 & \dots & 0 \\ 0 & \frac{2}{M} & \frac{M-2}{M} \frac{M-3}{M-2} & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & & \vdots \\ 0 & 0 & 0 & 0 & \frac{M-2}{M} & \frac{2}{M} \frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 & \frac{M-1}{M} \end{pmatrix}.$$

Thus, it is crucial for our analysis to understand the evolution of the Markov processes associated with matrices C_{\lfloor} and S_{\lfloor} , starting out with the situation in which only one player is infected. Then, for the buyer case, let $y_B^1 = (1, 0, \dots, 0)$ and define y_B^{t+1} as

$$y_B^{t+1} = \frac{y_B^t C_{\lfloor}}{\|y_B^t C_{\lfloor}\|} = \frac{y_B^1 C_{\lfloor}^t}{\|y_B^1 C_{\lfloor}^t\|}.$$

Analogously, we define the Markov process for the seller, y_S^t , by using S_{\lfloor} instead of C_{\lfloor} .

Therefore, my intermediate beliefs at the end of period \hat{T} , $x^{\hat{T}}$, would be given by $y_B^{\hat{T}}$ if I am a buyer and $y_S^{\hat{T}}$ if I am a seller. The following result characterizes the limit behavior of y_B^t and y_S^t , needed to understand $x^{\hat{T}}$ for large \hat{T} .

Lemma 1. Fix M . Then, $\lim_{t \rightarrow \infty} y_B^t = \lim_{t \rightarrow \infty} y_S^t = (0, \dots, 0, 1)$.

The intuition for the above lemma is as follows. Note that the largest diagonal entry in the matrix C_{\lfloor} (or S_{\lfloor}) is the last one. This means that the state $M - 1$ is more stable than

any other state. Consequently, as more periods of contagion elapse in Phase I, state $M - 1$ becomes more and more likely. The formal proof is a straightforward consequence of some of the properties of contagion matrices (see Proposition B.1 in the Online Appendix).

Proposition 2. *Fix \ddot{T} and M . If I observe history $h^{\dot{T}+\ddot{T}+1} = g \dots gb$ and \dot{T} is large enough, then it is sequentially rational for me to play the Nash action at period $\dot{T} + \ddot{T} + 2$.*

Proof. Suppose that I am a buyer. Since sellers always play the Nash action in Phase II, I cannot learn anything from play in Phase II. By Lemma 1, if \dot{T} is large enough, $x^{\dot{T}}$, which coincides with $y_B^{\dot{T}}$, will be close to $(0, \dots, 0, 1)$. Thus, I assign high probability to $M - 1$ players in my community being infected at the end of Phase I. Then, at least as many sellers got infected during Phase II. If exactly $M - 1$ buyers were infected by the end of Phase II, then the only uninfected seller must have got infected in period $\dot{T} + \ddot{T} + 1$ since, in this period, I was the only uninfected buyer and I met an infected opponent. So, if \dot{T} is large enough, I assign probability arbitrarily close to 1 to all sellers being infected by the end of $\dot{T} + \ddot{T} + 1$. Thus, it is optimal to play the Nash action a_2^* .

Next, suppose that I am a seller. In this case, the fact that no player infected me in Phase II will make me update my beliefs about how contagion has spread. However, if \dot{T} is large enough relative to \ddot{T} , even if I factor in the information that I was not infected in Phase II, the probability I assign to $M - 1$ buyers being infected by the end of Phase I is arbitrarily higher than the probability I assign to k players being infected for any $k < M - 1$. By the same argument as above, playing Nash is optimal. \square

Next, we consider $(\dot{T} + \ddot{T} + 1 + \alpha)$ -period histories of the form $h^{\dot{T}+\ddot{T}+1+\alpha} = g \dots gbg \dots^\alpha$, with $1 \leq \alpha \leq M - 2$, i.e., histories where I was infected at period $\dot{T} + \ddot{T} + 1$ and then I observed α periods of good behavior while I was playing the Nash action. For the sake of exposition, assume that I am a buyer (the arguments for a seller are analogous). Why are these histories significant? Notice that if I get infected in period $\dot{T} + \ddot{T} + 1$, I can believe that all other players in my community are infected. However, if after that I observe the on-path action, I will think that I have met an uninfected player and I will have to revise my beliefs, since it is not possible that all the players in my community were infected after period $\dot{T} + \ddot{T} + 1$. Can this alter my incentives to play the Nash action?

Suppose that $\alpha = 1$. After history $h^{\dot{T}+\ddot{T}+2} = g \dots gbg$, I know that at most $M - 2$ buyers were infected by the end of Phase I. So, for each $t \leq \dot{T}$, $x_M^t = x_{M-1}^t = 0$. My beliefs are no longer computed using $C_{1\downarrow}$, but rather with $C_{2\downarrow}$, which is derived by removing the last two rows and columns of C . By a similar argument as Lemma 1, we can show that

the new intermediate beliefs, $x^{\dot{T}} \in \mathbb{R}^{M-2}$, also converge to $(0, 0, \dots, 1)$. In other words, the state $M - 2$ is the most stable and, for \dot{T} large enough, I assign very high probability to $M - 2$ buyers being infected at the end of Phase I. Consequently, at least as many sellers were infected during Phase II. This, in turn, implies (just as in Proposition 2) that I believe that, with high probability, all players are infected by now (at $t = \dot{T} + \ddot{T} + 2$). To see why note that, with high probability, $M - 2$ sellers were infected during Phase II. Then, one of the uninfected sellers met an infected buyer at period $\dot{T} + \ddot{T} + 1$, and I infected the last one in the last period. So, if \dot{T} is large enough, I still assign arbitrarily high probability to everyone being infected by now, and it is optimal for me to play Nash. A similar argument holds for $(\dot{T} + \ddot{T} + 1 + \alpha)$ -period histories with $\alpha \in \{1, \dots, M - 1\}$.

We need not check incentives for Nash reversion after histories where I observe more than $M - 2$ periods of g after being infected. Note that these are erroneous histories. When I get infected at period $\dot{T} + \ddot{T} + 1$, I believe that the minimum number of people infected in each community is two (at least one pair of players in period 1 and one pair in period $\dot{T} + \ddot{T} + 1$). Suppose that subsequently I face $M - 2$ instances of g . Since I switched to Nash after being infected, I infected my opponents in the last $M - 2$ periods. At the end of $M - 2$ observations of g , I am sure that even if there were only two players infected in each community at the end of $\dot{T} + \ddot{T} + 1$, I have infected all the remaining uninfected people. So observing more g after that cannot be explained by a single deviation in period 1. Then, I will believe that in addition to the first deviation by a seller in period 1, there have been as many mistakes by players as needed to be consistent with the observed history.

Finally, consider histories where, after getting infected, I observe a sequence of actions that includes both g and b , *i.e.*, histories starting with $h^{\dot{T}+\ddot{T}+1} = g \dots gb$ and where I faced b in one or more periods after getting infected. After such histories, I will assign higher probability to more people being infected compared to histories where I observed only g after getting infected. This is because, by definition, every additional observation of b is an indication of more people being infected, which increases my incentives to play Nash.

Thus, we have shown that a player who observes a triggering action for the first time at the start of Phase III will find it optimal to revert permanently to the Nash action.

Case 2: Infection late in Phase III. Next, suppose that I get infected after observing history $h^{\bar{t}+1} = g \dots gb$, with $\bar{t} \gg \dot{T} + \ddot{T}$. This is the situation where the incentives are harder to verify, because we have to study how the contagion evolves during Phase III. In this phase, all infected players in both communities spread the contagion and, hence, from period $\dot{T} + \ddot{T} + 1$ on, the same number of people is infected in both communities. The

contagion can again be studied as a Markov process with state space $\{1, \dots, M\}$.

The new transition matrix corresponds with the contagion matrix $\bar{S} \in \mathcal{M}_M$ where, for each pair $k, l \in \{1, \dots, M\}$, if $k > l$ or $l > 2k$, $\bar{s}_{kl} = 0$; otherwise, *i.e.*, if $k \leq l \leq 2k$, the probability of transition to state k to state l is (see Figure 3):

$$\bar{s}_{kl} = \frac{\left(\binom{k}{l-k} \binom{M-k}{l-k} (l-k)! \right)^2 (2k-l)!(M-l)!}{M!} = \frac{(k!)^2 ((M-k)!)^2}{((l-k)!)^2 (2k-l)!(M-l)!M!}.$$

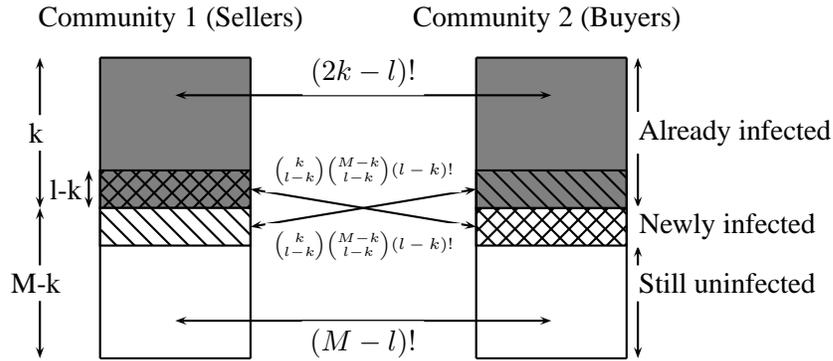


Figure 3: Spread of Contagion in Phase III. There are $M!$ possible matchings. For state k to transit to state l , exactly $(l-k)$ infected people from each community must meet $(l-k)$ uninfected people from the other community. The number of ways of choosing exactly $(l-k)$ buyers from k infected ones is $\binom{k}{l-k}$. The number of ways of choosing the corresponding $(l-k)$ uninfected sellers that will get infected is $\binom{M-k}{l-k}$. Finally, the number of ways in which these sets of $(l-k)$ people can be matched is the total number of permutations of $l-k$ people, *i.e.*, $(l-k)!$. Analogously, we choose the $(l-k)$ infected sellers who will be matched to $(l-k)$ uninfected buyers. The number of ways in which the remaining infected buyers and sellers get matched to each other is $(2k-l)!$ and, for the uninfected ones, we have $(M-l)!$.

Consider any t such that $\dot{T} + \ddot{T} < t < \bar{t}$. Since I have observed history $h^{\bar{t}+1} = g \dots gb$, I know that “at most $M-1$ people could have been infected in the rival community at the end of period $t+1$ ” ($\mathcal{I}^{t+1} \leq M-1$), and “I was not infected at the end of period $t+1$ ” (U^{t+1}). As before, let x^t be my intermediate beliefs after period t . Since, for each $t \leq \bar{t}$, $x_M^t = 0$, we can work with $x^t \in \mathbb{R}^{M-1}$. We want to compute $P(l^{t+1} | x^t \cap U^{t+1} \cap \mathcal{I}^{t+1} \leq M-1)$ for $l \in \{1, \dots, M-1\}$ and $\dot{T} + \ddot{T} < t$:

$$\begin{aligned} P(l^{t+1} | x^t \cap U^{t+1} \cap \mathcal{I}^{t+1} \leq M-1) &= \frac{P(l^{t+1} \cap U^{t+1} \cap \mathcal{I}^{t+1} \leq M-1) | x^t}{P(U^{t+1} \cap \mathcal{I}^{t+1} \leq M-1) | x^t} \\ &= \frac{\sum_{k \in \{1, \dots, M\}} x_k^t \bar{s}_{kl} \frac{M-l}{M-k}}{\sum_{l \in \{1, \dots, M-2\}} \left(\sum_{k \in \{1, \dots, M\}} x_k^t \bar{s}_{kl} \frac{M-l}{M-k} \right)}. \end{aligned}$$

Again, we can express these probabilities using the corresponding conditional transition matrix. Let $\bar{C} \in \mathcal{M}_M$ be defined, for each pair $k, l \in \{1, \dots, M-1\}$, by $\bar{c}_{kl} := \bar{s}_{kl} \frac{M-l}{M-k}$; by $\bar{c}_{MM} = 1$; and with all remaining entries being 0. Then, given a vector of beliefs at the end of Phase II represented by a probability vector \bar{y}^0 , we are interested in the evolution of the Markov process where \bar{y}^{t+1} is defined as

$$\bar{y}^{t+1} = \frac{\bar{y}^t \bar{C}_{\downarrow}}{\|\bar{y}^t \bar{C}_{\downarrow}\|}.$$

Note that, as long as $t \leq \bar{t} - \dot{T} - \ddot{T}$, each \bar{y}^t coincides with the intermediate beliefs $x^{\bar{T} + \dot{T} + t}$. Thus, to understand the beliefs after late deviations, we study \bar{y}^t as $t \rightarrow \infty$. Importantly, provided that $\bar{y}_1^0 > 0$, the limit results below do not depend on \bar{y}^0 , *i.e.*, the specific beliefs at the end of Phase II do not matter.³¹ We show next that a result similar to Lemma 1 holds.

Lemma 2. *Suppose that $\bar{y}_1^0 > 0$. Then, $\lim_{t \rightarrow \infty} \bar{y}^t = (0, 0, \dots, 0, 1) \in \mathbb{R}^{M-1}$.*

The lemma follows from properties of contagion matrices (see Proposition B.1 in the Online Appendix for the proof). The informal argument is as follows. The largest diagonal entries of the matrix \bar{C}_{\downarrow} are the first and last ones (\bar{c}_{11} and $\bar{c}_{M-1, M-1}$), which are equal. Unlike in the Phase I transition matrix, state $M-1$ is not the unique most stable state. Here, states 1 and $M-1$ are equally stable, and more stable than any other state. Why do then have converge to $(0, 0, \dots, 0, 1)$? In each period, many states transit to $M-1$ with positive probability, while no state transits to state 1, and so the ratio $\frac{\bar{y}_{M-1}^t}{\bar{y}_1^t}$ goes to ∞ as t increases.

Proposition 3. *Fix $\dot{T} \in \mathbb{N}$, $\ddot{T} \in \mathbb{N}$, and $M \in \mathbb{N}$. If I observe history $h^{\bar{t}+1} = g \dots gb$ and \bar{t} is large enough, then it is sequentially rational for me to play the Nash action at period $\bar{t} + 2$.*

Proof. We want to know how spread the contagion is according to the beliefs $x^{\bar{t}+1}$. To do so, we start with $x^{\bar{t}}$. By Lemma 2, if \bar{t} is large enough, $\bar{y}^{\bar{t} - \dot{T} - \ddot{T}}$ is very close to $(0, 0, \dots, 0, 1)$. Thus, $x^{\bar{t}}$ is such that I assign very high probability to $M-1$ players in the other community being infected by the end of period \bar{t} . Now, to compute $x^{\bar{t}+1}$ from $x^{\bar{t}}$, I add the information that I got infected at $\bar{t} + 1$ and, hence, the only uninfected person in the other community got infected too. Thus, I assign very high probability to everyone being infected, and Nash reversion is optimal. \square

³¹The condition $\bar{y}_1^0 > 0$ is just saying that, with positive probability, the deviant seller may meet the same buyer in all the periods in phases I and II.

Suppose now that I observe $h^{\bar{t}+2} = g \dots gbg$ and that I played the Nash action at period $\bar{t} + 2$. The analysis of these histories turns out to be more involved than in Case 1. After observing $h^{\bar{t}+2}$, I will know that fewer than $(M - 1)$ people were infected at the end of period \bar{t} since, otherwise, I could not have faced g in period $\bar{t} + 2$. In other words, I have to recompute my beliefs using the information that, for each $t \leq \bar{t}$, $\mathcal{I}^t \leq M - 2$. To do so, I should use the truncated transition matrix $\bar{C}_{2\downarrow}$. Since, for each $t \leq \bar{t}$, $x_M^t = x_{M-1}^t = 0$, to obtain $x^{\bar{t}}$ we just work with $x^t \in \mathbb{R}^{M-2}$. We have the Markov process that starts with a vector of beliefs at the end of Phase II, represented by a probability vector \bar{y}^0 , and such that $\bar{y}^{\bar{t}+1}$ is computed as

$$\bar{y}^{\bar{t}+1} = \frac{\bar{y}^{\bar{t}} \bar{C}_{2\downarrow}}{\|\bar{y}^{\bar{t}} \bar{C}_{2\downarrow}\|}.$$

As before, as long as $t \leq \bar{t} - \dot{T} - \ddot{T}$, each $\bar{y}^{\bar{t}}$ coincides with the intermediate beliefs $x^{\dot{T} + \ddot{T} + t}$. Also, we want to study the limit behavior of $\bar{y}^{\bar{t}}$ as t goes to ∞ and then rely on $x^{\bar{t}}$ to compute beliefs at $\bar{t} + 2$. The extra difficulty comes from the fact that $\bar{c}_{M-2, M-2} < \bar{c}_{11} = \frac{1}{M}$, and the arguments behind Lemma 2 do not apply to matrix $\bar{C}_{2\downarrow}$ and, indeed, the limit beliefs do not converge to $(0, \dots, 0, 1)$. However, we show that I will still believe that “enough” people are infected with “high enough probability.” In the lemma below, we first establish that $\bar{y}^{\bar{t}}$ converges. The result follows from properties of contagion matrices. The proof is available in the Online Appendix (Proposition B.1). Since this limit only depends on M , with a slight abuse of notation we denote it by \bar{y}^M .

Lemma 3. *Suppose that $\bar{y}_1^0 > 0$. Then, $\lim_{t \rightarrow \infty} \bar{y}^t = \bar{y}^M$, where \bar{y}^M is the unique nonnegative left eigenvector associated with the largest eigenvalue of $\bar{C}_{2\downarrow}$ such that $\|\bar{y}^M\| = 1$. Thus, $\bar{y}^M \bar{C}_{2\downarrow} = \frac{\bar{y}^M}{M}$.*

In particular, the result above implies that the limit as \bar{t} goes to infinity of the beliefs $x^{\bar{t}}$ is independent of \dot{T} and \ddot{T} . We are interested in the properties of \bar{y}^M , which we study next.

Lemma 4. *Let $r \in (0, 1)$. Then, for each $m \in \mathbb{N}$, there is $\underline{M} \in \mathbb{N}$ such that, for each $M \geq \underline{M}$,*

$$\sum_{j=\lceil rM \rceil}^{M-2} \bar{y}_j^M > 1 - \frac{1}{M^m},$$

where $\lceil z \rceil$ is the ceiling function and is defined as the smallest integer not smaller than z .

This result can be interpreted as follows. Think of r as a proportion of people, say 0.9. Given m , if the population size M is large enough, after observing history $h^{\bar{t}+2} = g \dots gbg$,

my limiting belief \bar{M} will be such that I will assign probability at least $(1 - \frac{1}{M^m})$ to at least 90 percent of the population being infected. We can choose r close enough to 1 and m large enough and then find $M \in \mathbb{N}$ large enough so that I believe that the contagion has spread enough that playing the Nash action is optimal.

Figure 4 below represents the probabilities $\sum_{j=\lceil rM \rceil}^{M-2} \bar{y}_i^M$ for different values of r and M . In particular, it shows that they go to one very fast with M .³² From the results in this section it will follow that, after any history in which I have been infected late in Phase III, I will believe that the contagion is at least as spread as \bar{y}^M indicates. For instance, consider $M = 20$. In this case, \bar{y}^M is such that I believe that, with probability at least 0.75, at least 90 percent of the people are infected. This may be enough to induce the right incentives for many games. In general, these incentives will hold even for fairly small population sizes.

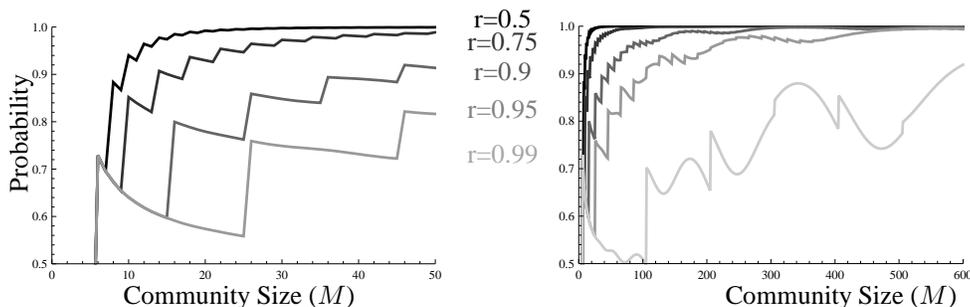


Figure 4: Illustration of Lemma 4.

In order to prove Lemma 4, we need to study the transitions between states. There are two opposing forces that affect how my beliefs evolve after I observe $g \dots gbg$. On the one hand, each observation of g suggests to me that not too many people are infected, making me step back in my beliefs and assign higher weight to lower states (fewer people infected). On the other hand, since I believe that the contagion started at $t = 1$ and that it has been spreading exponentially during Phase III, every elapsed period makes me assign more weight to higher states (believe that more people are infected). What we need to do is to compare the magnitudes of these two effects. Two main observations drive the proof. First, each time I observe g , my beliefs get updated with more weight assigned to lower states and, roughly speaking, this step back in beliefs turns out to be of the order of M . Second, we identify the most likely transition from any given state k . It turns out that the likeliest transition from k , say state k' , is about \sqrt{M} times more likely than the state k .

³²The non-monotonicity in the graphs in Figure 4 may be surprising. To the best of our understanding, this can be essentially attributed to the fact that the states that are powers of 2 tend to be more likely, and their distribution within the top $M - \lceil rM \rceil$ states varies in a non-monotone way.

Similarly, the most likely transition from k' , say state k'' , is \sqrt{M} times more likely than the state k' . Hence, given a proportion of people $r \in (0, 1)$ and $m \in \mathbb{N}$, if M is large enough, for each state $k < rM$, we can find another state \bar{k} that is at least M^m times more likely than state k . Thus, given $m \in \mathbb{N}$, by taking M large enough, we get that the second effect on the beliefs is, at least, of an order of M^m , and so it dominates the first effect.

We need some preliminaries before we formally prove Lemma 4. Recall that

$$(\bar{c}_2)_{k,k+j} = \bar{s}_{k,k+j} \frac{M-k}{M-k-j} = \frac{(k!)^2((M-k)!)^2}{(j!)^2(k-j)!(M-k-j)!M!} \frac{M-k-j}{M-k}.$$

Given a state $k \in \{1, \dots, M-2\}$, consider the transition from k to state $\text{tr}(k) := k + \lfloor \frac{k(M-k)}{M} \rfloor$, where $\lfloor z \rfloor$ is the floor function defined as the largest integer not larger than z . It turns out that, for large M , this is a good approximation of the most likely transition from state k .

We temporarily switch to the case where there is a continuum of states, *i.e.*, we think of the set of states as the interval $[0, M]$. In the continuous setting, a state $z \in [0, M]$, can be represented as rM ; where $r = z/M$ can be interpreted as the proportion of infected people at state z . Let $\gamma \in \mathbb{R}$. We define a function $f_\gamma: [0, 1] \rightarrow \mathbb{R}$ as

$$f_\gamma(r) := \frac{rM(M-rM)}{M} + \gamma = (r-r^2)M + \gamma.$$

Note that f_γ is continuous and, further, that $rM + f_0(r)$ would just be the extension of the function $\text{tr}(k)$ to the continuous case. We want to know how likely the transition from state rM to $rM + f_0(r)$ is. Define a function $g: [0, 1] \rightarrow [0, 1]$ as

$$g(r) := 2r - r^2.$$

The function g is continuous and strictly increasing. Given $r \in [0, 1]$, $g(r)$ represents the proportion of infected people if, at state rM , $f_0(r)$ people get infected. To see why, note that $rM + f_0(r) = rM + (r-r^2)M = (2r-r^2)M$. Let $g^2(r) := g(g(r))$ and define analogously any other power of g . Hence, for each $r \in [0, 1]$, $g^n(r)$ represents the fraction of people infected after n steps starting at rM when transitions are made according to $f_0(\cdot)$.

Claim 1. *Let $M \in \mathbb{N}$ and $a, b \in (0, 1)$, with $a > b$. Then, $aM + f_0(a) > bM + f_0(b)$.*

Proof. Note that $aM + f_0(a) - bM - f_0(b) = (g(a) - g(b))M$, and the result follows from the fact that $g(\cdot)$ is strictly increasing on $(0, 1)$. \square

Now we define a function $h_\gamma^M : (0, 1) \rightarrow (0, \infty)$ as

$$h_\gamma^M(r) := \frac{(rM!)^2((M - rM)!)^2}{(f_\gamma(r)!)^2(rM - f_\gamma(r))!(M - rM - f_\gamma(r))!M!} \frac{M - rM - f_\gamma(r)}{M - rM}.$$

This function is the continuous version of the transitions given by the matrix \bar{C}_2 . In particular, given $\gamma \in \mathbb{R}$ and $r \in [0, 1]$ the function $h_\gamma^M(r)$ represents the conditional probability of transition from state rM to state $rM + f_\gamma(r)$. In some abuse of notation, we apply the factorial function to non-integer real numbers. In such cases, the factorial can be interpreted as the corresponding Gamma function, *i.e.*, $a! = \Gamma(a + 1)$.

Claim 2. *Let $\gamma \in \mathbb{R}$ and $r \in (0, 1)$. Then, $\lim_{M \rightarrow \infty} Mh_\gamma^M(r) = \infty$. More precisely,*

$$\lim_{M \rightarrow \infty} \frac{Mh_\gamma^M(r)}{\sqrt{M}} = \frac{1}{r\sqrt{2\pi}}.$$

Proof. This proof is based on Stirling's approximation. The interested reader is referred to the Online Appendix (B.4) for a proof. \square

We are now ready to prove Lemma 4.

Proof of Lemma 4. Fix $r \in (0, 1)$ and $m \in \mathbb{N}$. We show that there is \underline{M} such that, for each $M \geq \underline{M}$ and each state $k < \lceil rM \rceil$, there is a state $\bar{k} \in \{\lceil rM \rceil, \dots, M - 2\}$ such that $\bar{y}_{\bar{k}}^M > M^{m+1}\bar{y}_k^M$. We show this first for state $k_0 = \lceil rM \rceil - 1$. Consider the state rM and define $\rho := 2m + 3$ and $\bar{r} := g^\rho(r)$. Recall that \bar{r} is the state reached (proportion of people infected) from initial state rM after ρ steps according to the function f_0 . Recall that functions f_0 and g are such that, $r < \bar{r} < 1$. Moreover, suppose that M is large enough so that $\bar{r}M \leq M - 2$. Consider the state k_0 and let \bar{k} be the number of infected people after ρ steps according to function $\text{tr}(\cdot)$. Given $\hat{r} \in (0, 1)$, $f_0(\hat{r})$ and $\text{tr}(\hat{r}M)$ just differ because $f_0(\hat{r})$ may not be a natural number, while $\text{tr}(\hat{r}M)$ always is, and so the difference between these two functions is just a matter of truncation. Hence, for each of step $j \in \{1, \dots, \rho\}$, there is $\gamma_j \in (-1, 0]$ such that the step corresponds with that of function f_{γ_j} ; for instance, γ_1 is such that $\text{tr}(k_0) = f_{\gamma_1}(k_0/M)$. By Claim 1, since $k_0 < rM$, $\bar{k} < M - 2$. Moreover, it is trivial to see that $\bar{k} > \lceil rM \rceil$. Let k_1 be the state that is reached after the first step from k_0 according to function $\text{tr}(\cdot)$. Recall that, by Lemma 3, $\bar{y}^M = M\bar{y}^M\bar{C}_2$. Then, $\bar{y}_{k_1}^M = M \sum_{k=1}^{M-2} \bar{y}_k^M (\bar{C}_2)_{kk_1} > M\bar{y}_{k_0}^M (\bar{C}_2)_{k_0k_1} = \bar{y}_{k_0}^M Mh_{\gamma_1}^M(r)$, which, by Claim 2, can be approximated by $\frac{\sqrt{M}}{r\sqrt{2\pi}}\bar{y}_{k_0}^M$ if M is large enough. Repeating the same argument for the other

intermediate states that are reached in each of the ρ steps we get that, if M is large enough,

$$\bar{y}_k^M > \frac{M^{\frac{\rho}{2}}}{(r\sqrt{2\pi})^\rho} \bar{y}_{k_0}^M = M^{m+1} \frac{M^{\frac{1}{2}}}{(r\sqrt{2\pi})^\rho} \bar{y}_{k_0}^M > M^{m+1} \bar{y}_{k_0}^M.$$

The proof for an arbitrary state $k < \lceil rM \rceil - 1$ is very similar, with the only difference that more steps might be needed to get to a state $\bar{k} \in \{\lceil rM \rceil, \dots, M - 2\}$; yet, the extra number of steps makes the difference between \bar{y}_k^M and $\bar{y}_{\bar{k}}^M$ even bigger.³³

Let $\bar{k} \in \{\lceil rM \rceil, \dots, M - 2\}$ be a state such that $\bar{y}_{\bar{k}}^M > M^{m+1} \max\{\bar{y}_k^M : k \in \{1, \dots, \lceil rM \rceil - 1\}\}$. Then, $\sum_{k=1}^{\lceil rM \rceil - 1} \bar{y}_k^M < rM \frac{\bar{y}_{\bar{k}}^M}{M^{m+1}} < \frac{1}{M^m}$. \square

Corollary 1. *Fix a game $G \in \mathcal{G}$. There is $M_G^1 \in \mathbb{N}$ such that, for each $M \geq M_G^1$, an infected player in Phase III with beliefs given by \bar{y}^M strictly prefers to play the Nash action.*

Proof. Suppose an infected player with beliefs given by \bar{y}^M is deciding what to do at some period in Phase III. By Lemma 4, the probability that he assigns to meeting an infected player in the current period converges to 1 as M goes to infinity. This has two implications on incentives. First, in the current period, the Nash action becomes a strict best reply. Second, the impact of the current action on the continuation payoff goes to 0. Therefore, there is $M_G^1 \in \mathbb{N}$ such that, for each M larger than M_G^1 , playing the Nash action is a best reply given \bar{y}^M . \square

Proposition 4. *Fix a game $G \in \mathcal{G}$ and $M > M_G^1$. Fix $\bar{T} \in \mathbb{N}$ and $\ddot{T} \in \mathbb{N}$ and let $\bar{t} \gg \bar{T} + \ddot{T}$. Suppose I observe history $h^{\bar{t}+2} = g \dots gbg$. Then, it is sequentially rational for me to play the Nash action at period $\bar{t} + 3$.*

Proof. Now, we want to know how spread the contagion is according to the beliefs $x^{\bar{t}+2}$. As usual, to do so we start with $x^{\bar{t}}$, my beliefs computed conditioning only the information that at most $M - 2$ people were infected after period \bar{t} and that I was uninfected until period $\bar{t} + 1$. From Lemma 3, if \bar{t} is large enough, $x^{\bar{t}}$, which coincides with $\bar{y}^{\bar{t}-\bar{T}-\ddot{T}}$, is very close to \bar{y}^M and, further, also all subsequent elements $\bar{y}^{\hat{t}-\bar{T}-\ddot{T}}$ with $\hat{t} > \bar{t}$ are very close to \bar{y}^M . In particular, we can assume that, taking $m = 1$ in Lemma 4, I believe that, with probability at least $1 - \frac{1}{M}$, at least rM people are infected. We can use $x^{\bar{t}}$ to compute my beliefs at the end of period $\bar{t} + 2$.

³³It is worth noting that we can do this argument uniformly for the different states and the corresponding γ 's because we know that all of them lie inside $[-1, 0]$, a bounded interval; that is, we can take M big enough so as to ensure that we can use the approximation given by Lemma 2 for all γ in $[-1, 0]$.

- After period $\bar{t} + 1$: I compute $x^{\bar{t}+1}$ by updating $x^{\bar{t}}$, conditioning on i) I observed b in period $\bar{t} + 1$ and ii) at most $M - 1$ people were infected after $\bar{t} + 1$ (I observed g at $\bar{t} + 2$). Let $\tilde{x}^{\bar{t}+1}$ be the belief computed from $x^{\bar{t}}$ by conditioning instead on i) I observed g in period $\bar{t} + 1$ and ii) at most $M - 2$ people are infected. Clearly, $x^{\bar{t}+1}$ first-order stochastically dominates $\tilde{x}^{\bar{t}+1}$, in the sense of placing higher probability on more people being infected. Moreover, $\tilde{x}^{\bar{t}+1}$ coincides with $\bar{y}^{\bar{t}+1-\bar{t}-\bar{t}}$, which we argued above is also close enough to \bar{y}^M .
- After period $\bar{t} + 2$: I compute $x^{\bar{t}+2}$ based on $x^{\bar{t}+1}$ and conditioning on i) I observed g ; ii) I infected my opponent by playing the Nash action at $\bar{t} + 2$; and iii) at most M people have been infected after $\bar{t} + 2$. Again, this updating leads to beliefs that first-order stochastically dominate $\tilde{x}^{\bar{t}+2}$, the beliefs we would obtain if we instead conditioned on i) I observed g and ii) at most $M - 2$ people were infected after $\bar{t} + 2$. Again, the belief $\tilde{x}^{\bar{t}+2}$ would be very close to \bar{y}^M .

Hence, if it is optimal for me to play the Nash action when my beliefs are given by \bar{y}^M , it is also optimal to do so after observing the history $h^{\bar{t}+2} = g \dots gbg$ (provided that \bar{t} is large enough). By Corollary 1, if $M > M_G^1$, the beliefs \bar{y}^M ensure that I have the incentive to switch to Nash. \square

As we did in Case 1, we still need to show that Nash reversion is optimal after histories of the form $h^{\bar{t}+1+\alpha} = g \dots gbg \dots g$. The intuition is as follows. If α is small, then I will think that a lot of players were already infected when I got infected, the argument being similar to that in Proposition 4. If α is large, I may learn that there were not so many players infected when I got infected. However, the number of players I myself have infected since then, together with the exponential spread of the contagion, will be enough to convince me that, at the present period, contagion is widespread anyway.

To formalize the above intuition, we need the following strengthening of Lemma 4. We omit the proof, as it just involves a minor elaboration of the arguments in Lemma 4.

Lemma 5. *Let $r \in (0, 1)$. Then, for each $m \in \mathbb{N}$, there are $\hat{r} \in (r, 1)$ and $\underline{M} \in \mathbb{N}$ such that, for each $M \geq \underline{M}$,*

$$\frac{\sum_{j=\lceil rM \rceil}^{\lceil \hat{r}M \rceil} \bar{y}_j^M}{1 - \sum_{j=\lfloor M-\hat{r}M \rfloor+1}^M \bar{y}_j^M} > 1 - \frac{1}{M^m}.$$

Proposition 5. Fix a game $G \in \mathcal{G}$. Then, there is M_G^2 such that, regardless of $\dot{T} \in \mathbb{N}$ and $\ddot{T} \in \mathbb{N}$, for each $\bar{t} \gg \dot{T} + \ddot{T}$ the following holds. For each $M > M_G^2$, if I observe history $h^{\bar{t}+1+\alpha} = g \dots gbg \dots^\alpha \cdot g$ with $\alpha \geq 1$, then it is sequentially rational for me to play the Nash action at period $\bar{t} + 2 + \alpha$.

Proof. The proof is similar to that of Proposition 4, so we relegate it to the Online Appendix (B.4). \square

Finally, we consider histories in which, after getting infected, I observe actions that include both g and b , i.e., histories starting with $h^{\bar{t}+1} = g \dots gbg$ and in which I have observed b in one or more periods after getting infected. The reasoning is the same as in Case 1. After such histories, I will assign higher probability to more people being infected compared to histories where I only observed g after getting infected.

Case 3: Infection in other periods of Phase III (“monotonicity” of beliefs). So far, we have shown that if a player is infected early in Phase III, he thinks that he was the last player to be infected and that everybody is infected, making Nash reversion optimal. Also, if a player is infected late in Phase III, he will believe that enough players are already infected for it to be optimal to play the Nash action (with his limit belief being given by \bar{y}^M). Next, we show that the belief of a player who is infected not very late in Phase III will lie in between. The earlier a player gets infected in Phase III, the closer his belief will be to $(0, \dots, 0, 1)$, and the later he gets infected, the closer his belief will be to \bar{y}^M .

Proposition 6. Fix a game $G \in \mathcal{G}$ and let $M > \max\{M_G^1, M_G^2\}$. Fix $\dot{T} \in \mathbb{N}$. If \dot{T} is large enough, then it is sequentially rational for me to play the Nash action after any history in which I get infected in Phase III.

Proof. Let $M > \max\{M_G^1, M_G^2\}$. In Cases 1 and 2, we showed that if I get infected at the start of Phase III (at $\dot{T} + \ddot{T} + 1$) or late in Phase III (at $\bar{t} \gg \dot{T} + \ddot{T}$), I will switch to the Nash action. What remains to be shown is that the same is true if I get infected at some intermediate period in Phase III. We prove this for histories in Phase III of the form $h^{\hat{t}+2} = g \dots gbg$. The proof can be extended to include other histories, just as in Cases 1 and 2. We want to compute my belief $x^{\hat{t}+2}$ after history $h^{\hat{t}+2} = g \dots gbg$. We first compute the intermediate belief x^t , for $t \leq \hat{t}$.

Beliefs are computed using matrix C_{2j} in Phase I, and \bar{C}_{2j} in Phase III. We know from Case 1 that, for \dot{T} large enough, $x^{\dot{T}+\ddot{T}+1} \in \mathbb{R}^M$ is close to $(0, \dots, 0, 1)$. Using the properties of the contagion matrix \bar{C}_{2j} (see Proposition B.2 in the Online Appendix), we can show that

if we start Phase III with such a belief $x^{\hat{T}+\hat{T}+1}$, $x^{\hat{t}}$ first-order stochastically dominates \bar{y}^M , in the sense of placing higher probability on more people being infected. My beliefs still need to be updated from $x^{\hat{t}}$ to $x^{\hat{t}+1}$ and then from $x^{\hat{t}+1}$ to $x^{\hat{t}+2}$. We can use similar arguments as in Proposition 4 to show that $x^{\hat{t}+2}$ first-order stochastically dominates \bar{y}^M . Hence, if it is sequentially rational for me to play the Nash action when my beliefs are given by \bar{y}^M , it is also sequentially rational to do so when they are given by $x^{\hat{t}+2}$. \square

We have established that if a player observes a triggering action during Phase III, it is sequentially rational for him to revert to the Nash action.

A.3.4 A player observes a triggering action in Phases I or II

It remains to check the incentives for a player who is infected during the trust-building phases. We argued informally in Section 3.2 that buyers infected in Phase I find it optimal to start the punishment only when Phase II starts. To get this, all that is needed is that, in each period of Phase I, the likelihood of meeting the deviant seller is low enough. For each game G , there is M_G^3 such that, if $M > M_G^3$, this will be the case. Also, we argued that in case of being infected in Phase II, they would find it optimal to switch to the Nash action. We omit the formal proofs, as the arguments are very similar to those for Case 1 above.

A.3.5 A player observes a non-triggering action

An uninfected player who observes a non-triggering action knows that his opponent will not get infected, and will continue to play as if on-path. Since he knows that contagion will not start, the best thing to do is also to ignore this off-path behavior.

A.3.6 Off-path histories with more than one deviation

First, consider the situation in which the deviating player observes a probability zero history. In such a case, he will be at an erroneous history (Assumption 3 in Section A.2) and will think that there have been some mistake, update his beliefs accordingly and best respond.

Second, suppose I am an infected player who deviates from the prescribed off-path action. The strategies prescribe that I subsequently play ignoring my own deviation (*i.e.*, play the Nash action). To see why this is optimal consider a history in which I have been infected at a period $\bar{t}+1$ late in Phase III and have observed history $h^{\bar{t}+1+\alpha} = g \dots gbg \dots g$.

Further, suppose that, instead of playing Nash, I have played my on-path action after being infected.

To argue why I will still think that contagion is widely spread regardless of how large α is,³⁴ let us compare the Markov processes needed to compute my intermediate beliefs with the one used to obtain the beliefs after history $h^{\bar{t}+1} = g \dots gb$.

History $h^{\bar{t}+1} = g \dots gb$. For this case we built upon the properties of \bar{y}^M to check the incentives (see Corollary 1). Recall that \bar{y}^M is the limit of the Markov process starting with some probability vector y^0 with $y_1^0 > 0$ and transition matrix $\bar{C}_{[2]}$.

History $h^{\bar{t}+1+\alpha} = g \dots gbg \dots g$. We start with the intermediate beliefs $x^{\bar{t}}$. Regardless of the value of α , since I am not spreading contagion (I may be meeting the same uninfected player in every period since I got infected), I will still think that at most $M - 2$ people were infected at any period $t \leq \bar{t}$. As above, the transition matrix is $\bar{C}_{[2]}$, and $x^{\bar{t}}$ will be close to \bar{y}^M . To compute subsequent beliefs, since I know that at least two people in each community were infected after \bar{t} , we have to use matrix $C_{[1,2]}$, which shifts the beliefs towards more people being infected (relative to the process given by $\bar{C}_{[2]}$). Therefore, the ensuing process will converge to a limit that stochastically dominates \bar{y}^M in terms of more people being infected.

Finally, to study the beliefs after histories in which, after being infected, I alternate on-path play with the Nash action, we would have to combine the above arguments with those in the proof of Proposition 5 (Online Appendix, B.4).

A.4 Choice of Parameters for Equilibrium Construction

We now describe the order in which the parameters, \underline{M} , \ddot{T} , \dot{T} , and δ , must be chosen to satisfy all incentive constraints simultaneously. Fix any game $G \in \mathcal{G}$ with a strict Nash equilibrium a^* and a target payoff $v \in F_{a^*}$.

The first step is to choose $\underline{M} > \max\{M_G^1, M_G^2, M_G^3\}$. This ensures that a buyer who gets infected in Phase I starts punishing in Phase II and that the incentive constraints in Phase III are satisfied, *i.e.*, a player who observes a triggering action late in Phase III believes that enough people are already infected so that Nash reversion is optimal. As we argued in the discussion after Lemma 4, $M_G^1 \approx 20$ should suffice for many games. A similar argument would apply to M_G^2 . Finally, M_G^3 will typically be smaller. However,

³⁴Note that the arguments in Proposition 5 cannot be applied since I am not spreading the contagion.

these are rough estimates. The precise computation of tight bounds is beyond the scope of this paper.

There is one subtle issue. Once \underline{M} , \ddot{T} , and \dot{T} are chosen, we need players to be patient enough to prevent deviations on-path. But we also need to check that a very patient player does not want to slow down the contagion once it has started. The essence of the argument is in observing that, for a fixed population size, once contagion has started, the expected future stage-game payoffs go down to the Nash payoff $u(a^*)$ exponentially fast. That is, regardless of how an infected player plays in the future, the undiscounted sum of possible future gains he can make relative to the payoff from Nash reversion is bounded above. Thus, in some sense, even a perfectly patient player becomes effectively impatient. In the Online Appendix (B.5), we show formally that even an infinitely patient player would have the incentives to revert to the Nash action when prescribed to do so by our strategies.

Once \underline{M} is chosen, we pick \ddot{T} large enough so that a buyer who is infected in Phase I and knows that not all buyers were infected by the end of Phase I still has an incentive to revert to Nash in Phase II. This buyer knows that contagion will spread from Phase III anyway, and the Nash action gives her a short-term gain in Phase II. Hence, if \ddot{T} is long enough, she will revert to Nash in Phase II.

Next, \dot{T} is chosen large enough so that i) a buyer infected in Phase I who has observed \hat{a}_1 in most periods of Phase I believes that with high probability all buyers were infected during Phase I; ii) a seller infected in Phase II believes that with high probability at least $M - 1$ buyers were infected during Phase I; iii) a seller who deviates in Phase I believes that, with high probability, he met all the buyers in Phase I; and iv) players infected in Phase III believe that, with high probability, “enough” people were infected by the end of Phase II.

Finally, once \underline{M} , \ddot{T} , and \dot{T} have been chosen, we find the threshold $\underline{\delta}$ such that i) for discount factors $\delta > \underline{\delta}$, players will not deviate on-path and ii) the payoff associated with our strategies is as close to the target payoff as needed.

B For Online Publication

B.1 Sequential Equilibrium: Consistency of Beliefs

In this section we prove that the beliefs we have defined are consistent. First, recall the assumptions used to characterize them.

- i) **Assumption 1:** If a player observes a triggering action, then this player believes that some player in Community 1 (a seller) deviated in the first period of the game. Subsequently, his beliefs are described as follows:
 - **Assumption 1.1:** He believes that after the deviation by a seller in the first period, play has proceeded as prescribed by the strategies, provided his observed history is can be explained in this way.
 - **Assumption 1.2:** If his observed history cannot be explained with play having proceeded as prescribed by the strategies after a seller's deviation in period 1, then we call this an *erroneous history*.³⁵ In this case, the player believes that after the seller deviated in the first period of the game, one or more of his opponents in his matches made a mistake. Indeed, this player will think that there have been as many mistakes by his past rivals as needed to explain the history at hand.
- ii) **Assumption 2:** If a player observes a non-triggering action, then this player believes that his opponent made a mistake.
- iii) **Assumption 3:** If a player plays a triggering action and then observes a history that has probability zero according to his beliefs, we also say that he is at an erroneous history and will think that there have been as many mistakes by his past rivals as needed to explain the history at hand.

Let σ and μ denote the equilibrium strategy profile and the associated system of beliefs, respectively. We prove below that there is a sequence of completely mixed strategies converging uniformly to σ such that the associated sequence of systems of beliefs, uniquely derived from the strategies by Bayes rule, converges pointwise to μ . Indeed, convergence will be uniform except for a set of pathological histories that has extremely low probability

³⁵For the formation of a player's beliefs after erroneous histories, we assume that "mistakes" by infected players are infinitely more likely than "mistakes" by uninfected players.

(and therefore, would have no influence at all on incentives). It is worth noting that there is hardly any reference in the literature in which consistency of beliefs is formally proved, so it is not even clear whether uniform convergence should be preferred over pointwise convergence. For the reader keen on former, at the end of this section we present a mild assumption on the matching technology under which uniform convergence to the limit beliefs would be obtained after every history.

Before moving to the proof, we present the notions of pointwise and uniform convergence adapted to our setting (we present them for beliefs, with the definitions for strategies being analogous):

Pointwise convergence. A sequence of beliefs $\{\mu_n\}_{n \in \mathbb{N}}$ converges pointwise to the belief μ if, for each $\delta > 0$ and each $T \in \mathbb{N}$, there is $\bar{n} \in \mathbb{N}$ such that, for each $n > \bar{n}$ and each information set at a period $t \leq T$, the probability attached by μ_n at each node of that information set differs by less than δ from that attached by μ .

Uniform convergence. A sequence of beliefs $\{\mu_n\}_{n \in \mathbb{N}}$ converges uniformly to the belief μ if, for each $\delta > 0$, there is $\bar{n} \in \mathbb{N}$ such that, for each $n > \bar{n}$ and each information set, the probability attached by μ_n at each node of that information set differs by less than δ from that attached by μ .

Proof of consistency. First, we construct the sequence of completely mixed strategies converging to our equilibrium strategies, σ , and let μ be a belief system satisfying assumptions 1-3. Fix a player i and let $D + 1$ be the number of actions available to player i in the stage game. Then, for each $n \in \mathbb{N}$, let $\varepsilon_n = \frac{1}{2n}$. Now, we distinguish two cases:

A buyer at period $t \leq \bar{T}$. An uninfected buyer plays the prescribed equilibrium action with probability $(1 - \varepsilon_n^t)$ and plays any other action with probability $\frac{\varepsilon_n^t}{D}$. For an infected buyer these probabilities are $(1 - \varepsilon_n^{nt})$ and $\frac{\varepsilon_n^{nt}}{D}$, respectively. These trembles convey the information that early deviations are more likely and that, in Phase I, mistakes by uninfected buyers are more likely than mistakes by infected ones.

A buyer at period $t > \bar{T}$ or a seller at any period. Each uninfected player plays the prescribed equilibrium action with probability $(1 - \varepsilon_n^{nt})$ and plays any other action with probability $\frac{\varepsilon_n^{nt}}{D}$. For an infected player these probabilities are $(1 - \varepsilon_n^t)$ and $\frac{\varepsilon_n^t}{D}$, respectively. These trembles convey again the information that early deviations are more likely, but now infected players are more likely to make mistakes.

Note that the above beliefs are natural in the spirit of trembling hand perfection: more costly mistakes are less likely. This is why except for the case of uninfected buyers in Phase I, whose mistakes are ignored anyway, mistakes by infected players are more likely (contagion will spread anyway).

For each $n \in \mathbb{N}$ we denote by σ_n the corresponding completely mixed strategy and by μ_n the belief system derived from σ_n using Bayes rule. Clearly, the sequence $\{\sigma_n\}_{n \in \mathbb{N}}$ converges uniformly to σ . We move now to the convergence of $\{\mu_n\}_{n \in \mathbb{N}}$ to μ , which we discuss through a series of representative histories. We start with the case that is more important for the incentives of our equilibrium strategies.

Case 1. Consider a history $h^{\bar{t}}$, late in Phase III ($\bar{t} \gg \dot{T} + \ddot{T}$), such that player i , a buyer, gets infected at time \bar{t} , i.e., $h^{\bar{t}} = g \dots gb$. We show that, as n goes to ∞ , the beliefs μ_n are such that the probability that player i attaches to a seller having deviated in period 1 converges to one.

For each $\bar{t} \in \mathbb{N}$ and $\tau \in \mathbb{N}$, with $\tau \leq \bar{t}$ we define $\mathcal{H}(\bar{t}, \tau)$ as the set of all histories of the matching process such that:

- Player i observed history $h^{\bar{t}} = g \dots gb$.
- The first triggering action occurred at period τ .

Let $\mathcal{P}_n(\bar{t}, \tau)$ denote the sum of the probabilities of the histories in $\mathcal{H}(\bar{t}, \tau)$ according to μ_n . Then, we want to show that, as n goes to ∞ , $\mathcal{P}_n(\bar{t}, 1)$ becomes infinitely larger than $\sum_{\tau=2}^{\bar{t}} \mathcal{P}_n(\bar{t}, \tau)$. To do so, we start comparing the probabilities in $\mathcal{H}(\bar{t}, 1)$ and $\mathcal{H}(\bar{t}, 2)$. Take $H \in \mathcal{H}(\bar{t}, 2)$ and associate with it an element $\hat{H} \in \mathcal{H}(\bar{t}, 1)$ such that:

- The same two players who got infected at period 2 according to H get infected in period 1 according to \hat{H} .
- These two players were matched to each other also in period 2.
- The realized matches in H and \hat{H} are the same from period 2 until \bar{t} .

We show that, as n goes to ∞ , player i assigns arbitrarily higher probability to \hat{H} compared to H . Recall that, at a given period t , the probability that a player plays the prescribed action is either $1 - \varepsilon_n^t$ or $1 - \varepsilon_n^{nt}$. Finally, note that triggering actions in Phase I must come from sellers and, therefore, all triggering actions have probability $\frac{\varepsilon_n^{n\tau}}{D}$. Then,

$$\frac{P(H)}{P(\hat{H})} \leq \frac{(1 - \varepsilon_n^n)^{2M} \frac{\varepsilon_n^{2n}}{D} (1 - \varepsilon_n^{2n})^{2M-1} q}{\frac{\varepsilon_n^n}{D} (1 - \varepsilon_n)^{2M-1} (1 - \varepsilon_n^2)^{2M} q},$$

where q is the probability of the realizations from period 2 until \bar{t} , which, by definition, is the same for H and \hat{H} . Simplifying the above expression we get

$$\varepsilon_n \frac{(1 - \varepsilon_n)^{2M} (1 - \varepsilon_n^{2n})^{2M-1}}{(1 - \varepsilon_n)^{2M-1} (1 - \varepsilon_n^2)^{2M}} \quad (1)$$

and, as n goes to ∞ , this converges to zero. We could have derived a similar expression if comparing histories in $\mathcal{H}(\bar{t}, \tau)$ and $\mathcal{H}(\bar{t}, \tau + 1)$. Further, \bar{t} would not appear in these expressions, which is important for our convergence arguments to hold uniformly (not just pointwise).

Next, to compare $\mathcal{P}_n(\bar{t}, 1)$ with $\mathcal{P}_n(\bar{t}, 2)$ we have to note the following. Given $H \in \mathcal{H}(\bar{t}, 2)$, there are $M!$ elements of $\mathcal{H}(\bar{t}, 2)$ that coincide with H from period 2 onwards (corresponding to all the possible matchings in period 1). On the other hand, given $\hat{H} \in \mathcal{H}(\bar{t}, 1)$, there are just $(M - 1)!$ elements of $\mathcal{H}(\bar{t}, 1)$ that coincide with \hat{H} from period 2 onwards (the players infected in period 2 must be matched also in period 1). What this implies is that if we want to compare $\mathcal{P}_n(\bar{t}, 1)$ and $\mathcal{P}_n(\bar{t}, 2)$ relying on the above association between elements of $\mathcal{H}(\bar{t}, 1)$ and $\mathcal{H}(\bar{t}, 2)$, we have to account for the fact that each element of $\mathcal{H}(\bar{t}, 1)$ will be associated with M different elements of $\mathcal{H}(\bar{t}, 2)$. Thus, if we define $\hat{\mathcal{H}}(\bar{t}, 1)$ as the set formed by the elements of $\mathcal{H}(\bar{t}, 1)$ that are associated with some element of $\mathcal{H}(\bar{t}, 2)$, we have that $|\mathcal{H}(\bar{t}, 2)| = M|\hat{\mathcal{H}}(\bar{t}, 1)|$. Then,

$$\frac{\mathcal{P}_n(\text{First dev. at } \tau = 2 | h^{\bar{t}})}{\mathcal{P}_n(\text{First dev. at } \tau = 1 | h^{\bar{t}})} = \frac{\mathcal{P}_n(\bar{t}, 2)}{\mathcal{P}_n(\bar{t}, 1)} = \frac{\sum_{H \in \mathcal{H}(\bar{t}, 2)} \mathcal{P}_n(H)}{\sum_{H' \in \mathcal{H}(\bar{t}, 1)} \mathcal{P}_n(H')} \leq \frac{\sum_{H \in \mathcal{H}(\bar{t}, 2)} \mathcal{P}_n(H)}{\sum_{\hat{H} \in \hat{\mathcal{H}}(\bar{t}, 1)} \mathcal{P}_n(\hat{H})}$$

which by the above argument and Equation 1 reduces to

$$M \varepsilon_n \frac{(1 - \varepsilon_n)^{2M} (1 - \varepsilon_n^{2n})^{2M-1}}{(1 - \varepsilon_n)^{2M-1} (1 - \varepsilon_n^2)^{2M}},$$

which, as we increase n , soon becomes smaller than ε_n^2 .

Wrapping up, we have shown that, according to the belief system μ_n , the event “the first triggering action occurred at period 1” is, for large n , at least $\frac{1}{\varepsilon_n^2}$ times more likely than the event “the first triggering action occurred at period 2”, with $\frac{1}{\varepsilon_n^2}$ converging to ∞ as n goes to ∞ .

More generally, the above arguments can be used, *mutandis mutandis*, to show that the event “the first triggering action occurred at period τ ” is, for large n , at least $\frac{1}{\varepsilon_n^2}$ times more

likely than the event “the first triggering action occurred at period $\tau + 1$ ”. Thus, in general we have that, for large n ,

$$\frac{\mathcal{P}_n(\bar{t}, \tau)}{\mathcal{P}_n(\bar{t}, 1)} \leq \frac{\varepsilon_n^2 \mathcal{P}_n(\bar{t}, \tau - 1)}{\mathcal{P}_n(\bar{t}, 1)} \leq \frac{\varepsilon_n^4 \mathcal{P}_n(\bar{t}, \tau - 2)}{\mathcal{P}_n(\bar{t}, 1)} \leq \dots \leq \frac{\varepsilon_n^{2(\tau-1)} \mathcal{P}_n(\bar{t}, 1)}{\mathcal{P}_n(\bar{t}, 1)} = \varepsilon_n^{2(\tau-1)}.$$

Therefore, for large n we have

$$\frac{\mathcal{P}_n(\text{First dev. at } \tau \neq 1 | h^{\bar{t}})}{\mathcal{P}_n(\text{First dev. at } \tau = 1 | h^{\bar{t}})} = \sum_{\tau=2}^{\bar{t}} \frac{\mathcal{P}_n(\bar{t}, \tau)}{\mathcal{P}_n(\bar{t}, 1)} \leq \sum_{\tau=2}^{\bar{t}} \varepsilon_n^{2(\tau-1)} \leq \sum_{\tau=2}^{\infty} \varepsilon_n^{2(\tau-1)} = \frac{\varepsilon_n^2}{1 - \varepsilon_n^2},$$

which, as we wanted to prove, converges to zero as n goes to ∞ and, further, this convergence is uniform on \bar{t} .

Case 2. Consider a history $h^{\bar{t}}$, in Phase I ($\bar{t} \leq \dot{T}$), such that player i , a seller, observes a nontriggering action at time \bar{t} , *i.e.*, $h^{\bar{t}} = g \dots gb$. In this case, what is relevant for the incentives is that the seller believes that there has been no triggering action so far.

First, by similar arguments to those in Case 1, we can show that, as n goes to ∞ , the set of histories with a triggering action in period 1 becomes infinitely more likely than the set of histories with the first triggering action at a period $\tau > 1$. We argue now why the former histories become themselves infinitely less likely than histories in which there has been no triggering action. The idea is that, to explain the observation of $h^{\bar{t}} = g \dots gb$ with a history that entails a triggering action at period 1, one needs at least two deviations from the equilibrium strategy: one at period 1, with probability $\frac{\varepsilon_n}{D}$ and another one by a buyer at period \bar{t} , which has probability at most $\frac{\varepsilon_n^{\bar{t}}}{D}$ (if the buyer was uninfected). On the other hand, explaining the observation of $h^{\bar{t}} = g \dots gb$ with histories without triggering actions just requires one deviation by a buyer at period \bar{t} which has probability $\frac{\varepsilon_n^{\bar{t}}}{D}$ (no one is infected). Comparing the likelihood of the two types of histories it is clear that the later ones become infinitely more likely as n goes to ∞ .

Case 3. We discuss now the beliefs after an erroneous history. Consider a history $h^{\bar{t}}$, at the start of Phase III ($\bar{t} = \dot{T} + \ddot{T} + 1$), such that player i , a buyer, gets infected at period \bar{t} and then observes $M - 1$ periods of good behavior while he is playing the Nash action: $h^{\bar{t}+M-1} = g \dots gbg \dots g$. As argued at the end of Case 1 in Section A.3.3, such a history is erroneous, since a deviation by a seller at period 1 does not suffice to explain it.³⁶ Then,

³⁶We briefly recall the argument used there. Player i believes that at least two people are infected at the end of period \bar{t} (at least one pair of players in period 1 and one pair in period \bar{t}). Suppose that subsequently

Assumption 1.2 on beliefs says that, in addition to the first deviation by a seller in period 1, there was another mistake; in this case, this means that one of the $M - 1$ opponents that player i faced after period \bar{t} was infected and, yet, he played as if on-path.

It is worth noting that history $h^{\bar{t}+M-1} = g \dots gbg \dots g$ can be alternatively explained by one deviation: the triggering action observed by player i at period \bar{t} was the first triggering action of the game and, hence, only one player in each community was infected after at the end of that period; in the other $M - 1$ subsequent periods player i infected each and every player in the other community one by one.

We informally argue why, as n goes to ∞ , histories like the one above become infinitely less likely than those with a seller deviating in period 1. The key is that deviations by infected players are much more likely than those by uninfected ones.

First triggering action at period 1. In this case the probability goes to zero at speed, at most, $\varepsilon_n^n \cdot \varepsilon_n^{\bar{t}+M-1} = \varepsilon_n^{n+\bar{t}+M-1}$.

First triggering action at period \bar{t} . As n goes to ∞ , the probability of these histories goes to zero at speed $\varepsilon_n^{n\bar{t}}$, which comes from the probability that an uninfected player deviates from the equilibrium path at period \bar{t} .

Clearly, as n goes to ∞ , $\varepsilon_n^{n\bar{t}}$ goes to zero faster than $\varepsilon_n^{n+\bar{t}+M-1}$. The same kind of comparison would arise for similar histories in which $\bar{t} \gg \dot{T} + \ddot{T}$, and therefore uniform convergence of beliefs would not be a problem either. Yet, a formal proof would require an analysis similar to that of Case 1, to factor in the likelihood of the matching realizations associated with different histories.

Case 4. To conclude, we discuss another type of erroneous histories, for which convergence is only achieved pointwise.³⁷ Consider a history in which player i makes the first triggering deviation of the game at period $\bar{t} = \dot{T} + \ddot{T} + 1$ and that, since then, he has always observed the Nash action $h = g \dots gbbb \dots$. Then, player i believes that, very likely, he has been meeting the same player, say player j , ever since he deviated (setting aside “mistakes”, no other player would have played the Nash action at period $\bar{t} + 1$, for instance). Given this belief suppose that, after deviating, player i played an action \tilde{a} that is

player i faces $M - 2$ instances of g . Then, since player i switched to the Nash action after getting infected, he has infected each and every one of his opponents in the last $M - 2$ periods. Then, at the end of $M - 2$ observations of g , he is sure that even if there were only two players infected in each community at the end of \bar{t} , he has infected all the remaining uninfected people. So observing more g after that cannot be explained by a single deviation in period 1.

³⁷We thank an anonymous referee for pointing this out.

neither the on-path action nor the Nash action (maybe maximizing short-run profits before the contagion spreads). The question now is, how does player j explain the history he is observing?

First of all, note that the types of histories we are describing are very pathological: players i and j are matched in every single period after period \bar{t} , which is very unlikely and leads to a situation in which, period after period, player i has very precise information about how the contagion is spreading. Since these histories have very low probability, they will have no impact on the incentives of a player who may be considering a deviation in Phase III but, in any case, we argue below why the beliefs of player j converge to those in which he thinks that a seller deviated in period 1 and all observations from period \bar{t} onwards have been mistakes. Let τ be the number of periods after \bar{t} in which players i and j have been matched and consider the following two histories:

First triggering action at period 1. The probability of these histories goes to zero at speed

$$\varepsilon_n^n \cdot \prod_{t=1}^{\tau} \varepsilon_n^{\bar{t}+t} = \varepsilon_n^n \cdot \varepsilon_n^{\tau \frac{2\bar{t}+\tau+1}{2}} = \varepsilon_n^{n+\tau \frac{2\bar{t}+\tau+1}{2}}$$

First triggering action at period \bar{t} . As n goes to ∞ , the probability of these histories goes to zero at speed $\varepsilon_n^{n\bar{t}}$, which comes from the probability that an uninfected player deviates from the equilibrium path.

Then, for each $\tau \in \mathbb{N}$, as n goes to ∞ , $\varepsilon_n^{n\bar{t}}$ goes to zero faster than $\varepsilon_n^{n+\tau \frac{2\bar{t}+\tau+1}{2}}$. However, differently from Case 3, the convergence cannot be made uniform. The reason is that for a fixed n , no matter how large, we can always find $\tau \in \mathbb{N}$ such that $\varepsilon_n^{n+\tau \frac{2\bar{t}+\tau+1}{2}}$ is smaller than $\varepsilon_n^{n\bar{t}}$; take, for instance, $\tau = n$. \square

B.1.1 Ensuring uniform convergence of beliefs

We have seen above that there are histories after which the convergence of beliefs required for consistency is achieved pointwise. These histories are very special, since they require the same two players being matched in many consecutive periods, which results in one of the players being highly informed about the evolution of contagion. As soon as there is a period in which these two players do not face each other, uniform convergence is recovered.

The preceding argument suggests a natural modification to our matching assumptions to ensure uniform convergence of beliefs after every history. The modification would be

to assume that there is $L \in \mathbb{N}$ such that no two players are matched to each other for more than L consecutive periods. In particular, $L = 1$ would mean that no players are matched together twice in a row. Thus, matching would not be completely independent over time. We conjecture that, regardless of the value of L , our construction goes through and convergence of beliefs would be uniform. However, a formal analysis of this modified setting is beyond the scope of this paper.

B.2 Properties of the Conditional Transition Matrices

In Section A.2, we introduced a class of matrices, contagion matrices, which turns out to be very useful in analyzing the beliefs of players. First, note that, since contagion matrices are upper triangular, their eigenvalues correspond with the diagonal entries. It is easy to check that, given a contagion matrix, any eigenvector associated with the largest eigenvalue is either nonnegative or nonpositive. Given $y \in \mathbb{R}^k$, let $\|y\| := \sum_{i \in \{1, \dots, k\}} y_i$. We are often interested in the limit behavior of $y^t := \frac{yC^t}{\|yC^t\|}$, where C is a contagion matrix and y is a probability vector. We present below a few results about this limit behavior. We distinguish three special types of contagion matrices that will deliver different limiting results.

Property C1: $\{c_{11}\} = \operatorname{argmax}_{i \in \{1, \dots, k\}} c_{ii}$.

Property C2: $c_{kk} \in \operatorname{argmax}_{i \in \{1, \dots, k\}} c_{ii}$.

Property C3: For each $l < k$, $C_{[l]}$ satisfies C1 or C2.

Lemma B.1. *Let C be a contagion matrix and let λ be its largest eigenvalue. Then, the left eigenspace associated with λ has dimension one. That is, the geometric multiplicity of λ is one, irrespective of its algebraic multiplicity.*

Proof. Let l be the largest index such that $c_{ll} = \lambda > 0$, and let y be a nonnegative left eigenvector associated with λ . We claim that, for each $i < l$, $y_i = 0$. Suppose not and let i be the largest index smaller than l such that $y_i \neq 0$. If $i < l - 1$, we have that $y_{i+1} = 0$ and, since $c_{i,i+1} > 0$, we get $(yC)_{i+1} > 0$, which contradicts that y is an eigenvector associated with λ . If $i = l - 1$, then $(yC)_l \geq c_{ll}y_l + c_{l-1,l}y_{l-1} > c_{ll}y_l = \lambda y_l$, which, again, contradicts that y is an eigenvector associated with λ . Then, we can restrict attention to matrix $C_{[(l-1)]}$. Now, λ is also the largest eigenvalue of $C_{[(l-1)]}$ but, by definition of l , only one diagonal entry of $C_{[(l-1)]}$ equals λ and, hence, its multiplicity is one. Then, $z \in \mathbb{R}^{k-(l-1)}$ is a left

eigenvector associated with λ for matrix $C_{[(l-1)}$ if and only if $(0, \dots, 0, z) \in \mathbb{R}^k$ is a left eigenvector associated with λ for matrix C . \square

Given a contagion matrix C with largest eigenvalue λ , we denote by \check{y} the unique non-negative left eigenvector associated with λ such that $\|\check{y}\| = 1$.

Proposition B.1. *Let $C \in \mathcal{M}_k$ be a contagion matrix satisfying C1 or C2. Then, for each nonnegative vector $y \in \mathbb{R}^k$ with $y_1 > 0$, we have $\lim_{t \rightarrow \infty} \frac{yC^t}{\|yC^t\|} = \check{y}$. In particular, under C2, $\check{y} = (0, \dots, 0, 1)$.*

Proof. Clearly, since C is a contagion matrix, if t is large enough, all the components of y^t are positive. Then, for the sake of exposition, we assume that all the components of y are positive. We distinguish two cases.

C satisfies C1. This part of the proof is a direct application of Perron-Frobenius theorem. First, note that $\frac{yC^t}{\|yC^t\|}$ can be written as $\frac{y(C^t/\lambda^t)}{\|y(C^t/\lambda^t)\|}$. Now, using for instance Theorem 1.2 in Seneta (2006), we have that $\frac{C^t}{\lambda^t}$ converges to a matrix that is obtained as the product of the right and left eigenvectors associated to λ . Since in our case the right eigenvector is $(1, 0, \dots, 0)$, $\frac{C^t}{\lambda^t}$ converges to a matrix that has \check{y} in the first row and with all other rows being the zero vector. Therefore, the result follows from the fact that $y_1 > 0$.

C satisfies C2. We show that, for each $i < k$, $\lim_{t \rightarrow \infty} y_i^t = 0$. We prove this by induction on i . Let $i = 1$. Then, for each $t \in \mathbb{N}$,

$$\frac{y_1^{t+1}}{y_k^{t+1}} = \frac{c_{11}y_1^t}{\sum_{l \leq k} c_{lk}y_l^t} < \frac{c_{11}y_1^t}{c_{kk}y_k^t} \leq \frac{y_1^t}{y_k^t},$$

where the first inequality is strict because $y_{k-1} > 0$ and $c_{k-1,k} > 0$ (C is a contagion matrix); the second inequality follows from C2. Hence, the ratio $\frac{y_1^t}{y_k^t}$ is strictly decreasing in t . Moreover, since all the components of y^t lie in $[0, 1]$, it is not hard to see that, as far as y_1^t is bounded away from 0, the speed at which the above ratio decreases is also bounded away from 0.³⁸ Therefore, $\lim_{t \rightarrow \infty} y_1^t = 0$. Suppose that the claim holds for each $i < j < k - 1$. Now,

$$\frac{y_j^{t+1}}{y_k^{t+1}} = \frac{\sum_{l \leq j} c_{lj}y_l^t}{\sum_{l \leq k} c_{lk}y_l^t} < \frac{\sum_{l \leq j} c_{lj}y_l^t}{c_{kk}y_k^t} = \sum_{l < j} \frac{c_{lj}}{c_{kk}} \frac{y_l^t}{y_k^t} + \frac{c_{jj}}{c_{kk}} \frac{y_j^t}{y_k^t} \leq \sum_{l < j} \frac{c_{lj}}{c_{kk}} \frac{y_l^t}{y_k^t} + \frac{y_j^t}{y_k^t}.$$

³⁸Roughly speaking, this is because the state k will always get some probability from state 1 via the intermediate states, and this probability will be bounded away from 0 as far as the probability of state 1 is bounded away from 0.

By the induction hypothesis, for each $l < j$, the term $\frac{y_l^t}{y_k^t}$ can be made arbitrarily small for large enough t . Then, the first term in the above expression can be made arbitrarily small. Hence, it is easy to see that, for large enough t , the ratio $\frac{y_j^t}{y_k^t}$ is strictly decreasing in t . As above, this can happen only if $\lim_{t \rightarrow \infty} y_j^t = 0$. \square

Recall the matrices used to represent a player's beliefs after he observes history $h^t = g \dots gb$. At the beginning of Phase III, the beliefs evolved according to matrices C_{1j} and S_{1j} , and late in Phase III, according to \bar{C}_{1j} . Note that these three matrices all satisfy the conditions of the above proposition. This is what drives Lemmas 1 and 2 in the text. Consider the truncated matrix \bar{C}_{2j} that gave the transition of beliefs of a player who observes history $h^t = g \dots bg$. This matrix also satisfies the conditions of the above proposition, and this suffices for Lemma 3.

Proposition B.2. *Let $C \in \mathcal{M}_k$ be a contagion matrix satisfying C1 and C3. Let $y \in \mathbb{R}^k$ be a nonnegative vector. Then, if y is close enough to $(0, \dots, 0, 1)$, we have that, for each $t \in \mathbb{N}$ and each $l \in \{1, \dots, k\}$, $\sum_{i=l}^k y_i^t \geq \sum_{i=l}^k \check{y}_i$.*

Whenever two vectors are as y^t and \check{y} above, we say that y^t *first-order stochastically dominates* \check{y} (in the sense of more people being infected).

Proof. For each $i \in \{1 \dots, k\}$, let e_i denote the i -th element of the canonical basis in \mathbb{R}^k . By C1, c_{11} is larger than any other diagonal entry of C . Let \check{y} be the unique nonnegative left eigenvector associated with c_{11} such that $\|\check{y}\| = 1$. Clearly, $\check{y}_1 > 0$ and, hence, $\{\check{y}, e_2, \dots, e_k\}$ is a basis in \mathbb{R}^k . With respect to this basis, the matrix C looks like

$$\left(\begin{array}{c|c} c_{11} & 0 \\ \hline 0 & C_{\bar{1}} \end{array} \right).$$

Now, we distinguish two cases.

$C_{\bar{1}}$ **satisfies C2.** In this case, we can apply Proposition B.1 to $C_{\bar{1}}$ to get that, for each nonnegative vector $z \in \mathbb{R}^{k-1}$ with $z_1 > 0$, $\lim_{t \rightarrow \infty} \frac{z C_{\bar{1}}^t}{\|z C_{\bar{1}}^t\|} = (0, \dots, 0, 1)$. Now, let $y \in \mathbb{R}^k$ be the vector in the statement of this result. Since y is very close to $(0, \dots, 0, 1)$. Then, using the above basis, it is clear that $y = \alpha \check{y} + v$, with $\alpha \geq 0$ and $v \approx (0, \dots, 0, 1)$. Let

$t \in \mathbb{N}$. Then, for each $t \in \mathbb{N}$,

$$y^t = \frac{yC^t}{\|yC^t\|} = \frac{\lambda^t \alpha \check{y} + vC^t}{\|yC^t\|} = \frac{\lambda^t \alpha \check{y} + \|vC^t\| \frac{vC^t}{\|vC^t\|}}{\|yC^t\|}.$$

Clearly, $\|yC^t\| = \|\lambda^t \alpha \check{y} + \|vC^t\| \frac{vC^t}{\|vC^t\|}\|$ and, since all the terms are positive,

$$\|yC^t\| = \|\lambda^t \alpha\| \|\check{y}\| + \|vC^t\| \left\| \frac{vC^t}{\|vC^t\|} \right\| = \|\lambda^t \alpha\| + \|vC^t\|$$

and, hence, we have that y^t is a convex combination of \check{y} and $\frac{vC^t}{\|vC^t\|}$. Since $v \approx (0, \dots, 0, 1)$ and $\frac{vC^t}{\|vC^t\|} \rightarrow (0, \dots, 0, 1)$, it is clear that, for each $t \in \mathbb{N}$, $\frac{vC^t}{\|vC^t\|}$ first-order stochastically dominates \check{y} in the sense of more people being infected. Therefore, also y^t will first-order stochastically dominate \check{y} .

$C_{\bar{1}}$ satisfies **C1**. By C1, the first diagonal entry of $C_{\bar{1}}$ is larger than any other diagonal entry. Let \check{y}^2 be the unique associated nonnegative left eigenvector such that $\|\check{y}^2\| = 1$. It is easy to see that \check{y}^2 first-order stochastically dominates \check{y} ; the reason is that \check{y}^2 and \check{y} are the limit of the same contagion process, with the only difference that the state in which only one person has been infected is known to have probability 0 when using obtaining \check{y}^2 from $C_{\bar{1}}$. Clearly, $\check{y}_2^2 > 0$ and, hence, $\{\check{y}, \check{y}^2, e_3, \dots, e_k\}$ is a basis in \mathbb{R}^k . With respect to this basis, the matrix C looks like

$$\left(\begin{array}{c|c|c} c_{11} & 0 & 0 \\ \hline 0 & c_{22} & 0 \\ \hline 0 & 0 & C_{\bar{2}} \end{array} \right).$$

Again, we can distinguish two cases.

- $C_{\bar{2}}$ satisfies **C2**. In this case, we can repeat the arguments above to show that y^t is a convex combination of \check{y} , \check{y}^2 and $\frac{vC^t}{\|vC^t\|}$. Since both \check{y}^2 and $\frac{vC^t}{\|vC^t\|}$ first-order stochastically dominate \check{y} , y^t also does.
- $C_{\bar{2}}$ satisfies **C1**. Now, we would get a vector \check{y}^3 , and the procedure would continue until a truncated matrix satisfies C2 or until we get a basis of eigenvectors, one of them being \check{y} and all the others first-order stochastically dominating \check{y} . In both

situations, the result immediately follows from the above arguments. \square

Note that the matrix \bar{C}_2 , which gave the transition of beliefs of a player conditional on history $h^t = g \dots gbg$ late in the game, satisfies the conditions of the above proposition. This property is useful in proving Proposition 6.

B.3 Updating of Beliefs Conditional on Observed Histories

Below, we validate the approach to computing off-path beliefs discussed in Section A.3.1. Suppose that player i observes history $h^{\bar{t}+1} = g \dots gbg$ in Phase III, and we want to compute her beliefs at period $\bar{t} + 1$ conditional on $h^{\bar{t}+1}$, namely $x^{\bar{t}+1}$. Recall our method for computing $x^{\bar{t}+1}$. We first compute a set of intermediate beliefs x^t for $t < \bar{t} + 1$. For any period τ , we compute $x^{\tau+1}$ from x^τ by conditioning on the event that “I was uninfected in period $\tau + 1$ ” and that “ $\mathcal{I}^{\tau+1} \leq M - 2$ ” (\mathcal{I}^t is the random variable representing the number of infected people after period t). We do not use the information that “I remained uninfected after any period t with $\tau + 1 < t < \bar{t}$.” This information is added later, period by period, *i.e.*, only at period t , we add the information coming from the fact that “I was not infected at period t .” Below, we show that this method of computing beliefs is equivalent to the standard updating of beliefs conditioning on the entire history at once.

Let $\alpha \in \{0, \dots, M - 2\}$ and let $h^{t+1+\alpha}$ denote the $(t+1+\alpha)$ -period history $g \dots gbg \dots^\alpha g$. Recall that U^t denotes the event that i is uninfected at the end of period t . Let $b_t(g_t)$ denote the event that player i faced $b(g)$ in period t . We introduce some additional notation.

- $I_{i,k}^{(t)}$ denotes the event $i < \mathcal{I}^t < k$, *i.e.*, the number of infected people at the end of t periods is at least i and at most k .
- $E_\alpha^t := I_{0, M-\alpha}^t \cap U^t$
- $E_\alpha^{t+1} := E_\alpha^t \cap I_{1, M-\alpha+1}^{t+1} \cap b^{t+1}$
- For each $\beta \in \{1, \dots, \alpha - 1\}$,

$$E_\alpha^{t+1+\beta} := E_\alpha^{t+\beta} \cap I_{\beta+1, M-\alpha+\beta+1}^{t+1+\beta} \cap g^{t+1+\beta}$$
- $E_\alpha^{t+1+\alpha} := E_\alpha^{t+\alpha} \cap g^{t+1+\alpha} = h^{t+1+\alpha}$.

Let H^t be a history of the contagion process up to period t . Let \mathcal{H}^t be the set of all H^t histories. \mathcal{H}_k^t denotes the set of t -period histories of the stochastic process where $\mathcal{I}^t = k$.

We say $H^{t+1} \Rightarrow h^{t+1}$ if history H^{t+1} implies that I observed history h^{t+1} . Let $i \xrightarrow{t+1+\beta} k$ denote the event that state i transits to state k in period $t+1+\beta$, consistently with $h^{t+1+\alpha}$, or, equivalently, consistently with $E_\alpha^{t+1+\beta}$.

The probabilities of interest are $P(\mathcal{I}^{t+1+\alpha} = k |_{h^{t+1+\alpha}}) = P(\mathcal{I}^{t+1+\alpha} = k |_{E_\alpha^{t+1+\alpha}})$. We want to show that we can obtain the probabilities after $t+1+\alpha$ conditional on $h^{t+1+\alpha}$ by starting with the probabilities after t conditional on E_α^t and then let the contagion elapse one more period at a time conditioning on the new information, *i.e.*, adding the “local” information that player i observed g in the next period and that infected one more person. Precisely, we want to show that, for each $\beta \in \{0, \dots, \alpha\}$,

$$P(\mathcal{I}^{t+1+\beta} = k |_{E_\alpha^{t+1+\beta}}) \stackrel{?}{=} \frac{\sum_{i=1}^M P(i \xrightarrow{t+1+\beta} k) P(\mathcal{I}^{t+\beta} = i |_{E_\alpha^{t+\beta}})}{\sum_{j=1}^M \sum_{i=1}^M P(i \xrightarrow{t+1+\beta} j) P(\mathcal{I}^{t+\beta} = i |_{E_\alpha^{t+\beta}})}.$$

Fix $\beta \in \{0, \dots, \alpha\}$. For each $H^{t+1+\beta} \in \mathcal{H}^{t+1+\beta}$, let $H^{t+1+\beta, \beta}$ denote the unique $H^{t+\beta} \in \mathcal{H}^{t+\beta}$ that is compatible with $H^{t+1+\beta}$, *i.e.*, the restriction of $H^{t+1+\beta}$ to the first $t+\beta$ periods. Let $F^{1+\beta} := \{\tilde{H}^{t+1+\beta} \in \mathcal{H}^{t+1+\beta} : \tilde{H}^{t+1+\beta} \Rightarrow E_\alpha^{t+1+\beta}\}$. Let $F_k^{1+\beta} := \{\tilde{H}^{t+1+\beta} \in F^{1+\beta} : \tilde{H}^{t+1+\beta} \in \mathcal{H}_k^{t+1+\beta}\}$. Clearly, the $F_k^{1+\beta}$ sets define a “partition” of $F^{1+\beta}$ (one or more sets in the partition might be empty). Let $F_k^\beta := \{\tilde{H}^{t+1+\beta} \in F^{1+\beta} : \tilde{H}^{t+1+\beta, \beta} \in \mathcal{H}_k^{t+\beta}\}$. Clearly, also the F_k^β sets define a “partition” of $F^{1+\beta}$. Note that, for each pair $H^{t+1+\beta}, \tilde{H}^{t+1+\beta} \in F_k^{1+\beta} \cap F_i^\beta$, $P(H^{t+1+\beta} |_{H^{t+1, \beta}}) = P(\tilde{H}^{t+1+\beta} |_{\tilde{H}^{t+1, \beta}})$. Denote this probability by $P(F_i^\beta \xrightarrow{t+1+\beta} F_k^{1+\beta})$. Let $|i \xrightarrow{t+1+\beta} k|$ denote the number of ways in which i can transition to k at period $t+1+\beta$ consistently with $h^{t+1+\alpha}$ or, equivalently, consistently with $E_\alpha^{t+1+\beta}$. Clearly, this number is independent of the history that led to i people being infected. Now, $P(i \xrightarrow{t+1+\beta} k) = P(F_i^\beta \xrightarrow{t+1+\beta} F_k^{1+\beta}) |i \xrightarrow{t+1+\beta} k|$. Then,

$$\begin{aligned} & P(\mathcal{I}^{t+1+\beta} = k |_{E_\alpha^{t+1+\beta}}) = \\ &= \sum_{H^{t+1+\beta} \in \mathcal{H}_k^{t+1+\beta}} P(H^{t+1+\beta} |_{E_\alpha^{t+1+\beta}}) = \sum_{H^{t+1+\beta} \in F_k^{1+\beta}} P(H^{t+1+\beta} |_{E_\alpha^{t+1+\beta}}) \\ &= \sum_{H^{t+1+\beta} \in F_k^{1+\beta}} \frac{P(H^{t+1+\beta} \cap E_\alpha^{t+1+\beta})}{P(E_\alpha^{t+1+\beta})} = \frac{1}{P(E_\alpha^{t+1+\beta})} \sum_{H^{t+1+\beta} \in F_k^{1+\beta}} P(H^{t+1+\beta}) \\ &= \frac{1}{P(E_\alpha^{t+1+\beta})} \sum_{i=1}^M \sum_{H^{t+1+\beta} \in F_k^{1+\beta} \cap F_i^\beta} P(H^{t+1+\beta}) \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{P(E_\alpha^{t+1+\beta})} \sum_{i=1}^M \sum_{H^{t+1+\beta} \in F_k^{1+\beta} \cap F_i^\beta} P(H^{t+1+\beta} |_{H^{t+1+\beta, \beta}}) P(H^{t+1+\beta, \beta} |_{E_\alpha^{t+\beta}}) P(E_\alpha^{t+\beta}) \\
&= \frac{P(E_\alpha^{t+\beta})}{P(E_\alpha^{t+1+\beta})} \sum_{i=1}^M P(F_i^\beta \xrightarrow{t+1+\beta} F_k^{1+\beta}) \sum_{H^{t+1+\beta} \in F_k^{1+\beta} \cap F_i^\beta} P(H^{t+1+\beta, \beta} |_{E_\alpha^{t+\beta}}) \\
&= \frac{P(E_\alpha^{t+\beta})}{P(E_\alpha^{t+1+\beta})} \sum_{i=1}^M P(F_i^\beta \xrightarrow{t+1+\beta} F_k^{1+\beta}) |i \xrightarrow{t+1+\beta} k| \sum_{H^{t+\beta} \in \mathcal{H}_i^{t+\beta}} P(H^{t+\beta} |_{E_\alpha^{t+\beta}}) \\
&= \frac{P(E_\alpha^{t+\beta})}{P(E_\alpha^{t+1+\beta})} \sum_{i=1}^M P(F_i^\beta \xrightarrow{t+1+\beta} F_k^{1+\beta}) |i \xrightarrow{t+1+\beta} k| P(\mathcal{I}^{t+\beta} = i |_{E_\alpha^{t+\beta}}) \\
&= \frac{P(E_\alpha^{t+\beta})}{P(E_\alpha^{t+1+\beta})} \sum_{i=1}^M P(i \xrightarrow{t+1+\beta} k) P(\mathcal{I}^{t+\beta} = i |_{E_\alpha^{t+\beta}})
\end{aligned}$$

It is easy to see that $P(E_\alpha^{t+1+\beta}) = \sum_{j=1}^M P(E_\alpha^{t+\beta}) \sum_{i=1}^M P(i \xrightarrow{t+1+\beta} j) P(\mathcal{I}^{t+\beta} = i |_{E_\alpha^{t+\beta}})$, and the result follows. Similar arguments apply to histories $h^{t+1+\alpha} = g \dots gbg \dots^\alpha$, where player i observes both g and b in the α periods following the first triggering action.

B.4 Proofs of Claim 2 and Proposition 5

Proof of Claim 2. We prove the claim in two steps.

Step 1: $\gamma = 0$. Stirling's formula implies that $\lim_{n \rightarrow \infty} (e^{-n} n^{n+\frac{1}{2}} \sqrt{2\pi}) / n! = 1$. Given $r \in (0, 1)$, to study $h_\gamma^M(r)$ in the limit, we use the approximation $n! = e^{-n} n^{n+\frac{1}{2}} \sqrt{2\pi}$. Substituting and simplifying, we get the following:

$$\begin{aligned}
Mh_0^M(r) &= M \frac{((rM)!)^2 (((1-r)M)!)^2}{M!(r^2M)!((r-r^2)M)!^2((1-r)^2M)!} (1-r) \\
&= \frac{M(rM)^{1+2rM} ((1-r)M)^{1+2(1-r)M} (1-r)}{\sqrt{2\pi} M^{\frac{1}{2}+M} ((1-r)^2M)^{1+2(1-r)^2M} ((r-r^2)M)^{\frac{1}{2}+(r-r^2)M} (r^2M)^{\frac{1}{2}+r^2M}} \\
&= \frac{\sqrt{M}}{r\sqrt{2\pi}}.
\end{aligned}$$

Step 2: Let $\gamma \in \mathbb{R}$ and $r \in (0, 1)$. Now,

$$\frac{h_0^M(r)}{h_\gamma^M(r)} = \frac{(r^2M - \gamma)! (((r-r^2)M + \gamma)!)^2 ((1-r)^2M - \gamma)!}{(r^2M)! (((r-r^2)M)!)^2 ((1-r)^2M)!} \frac{(1-r)^2M}{(1-r)^2M - \gamma}.$$

Applying Stirling's formula again, the above expression becomes

$$\frac{(r^2M - \gamma)^{\frac{1}{2} + r^2M - \gamma}}{(r^2M)^{\frac{1}{2} + r^2M}} \frac{((r - r^2)M + \gamma)^{1 + 2(r - r^2)M + 2\gamma}}{((r - r^2)M)^{1 + 2(r - r^2)M}} \frac{((1 - r)^2M - \gamma)^{\frac{1}{2} + (1 - r)^2M - \gamma}}{((1 - r)^2M)^{\frac{1}{2} + (1 - r)^2M - \gamma}} \frac{(1 - r)^{2M}}{(1 - r)^{2M - \gamma}}. \quad (2)$$

To compute the limit of the above expression as $M \rightarrow \infty$, we analyze the four fractions above separately. Clearly, $((1 - r)^2M)/((1 - r)^2M - \gamma) \rightarrow 1$ as $M \rightarrow \infty$. So, we restrict attention to the first three fractions. Take the first one:

$$\frac{(r^2M - \gamma)^{\frac{1}{2} + r^2M - \gamma}}{(r^2M)^{\frac{1}{2} + r^2M}} = \left(1 - \frac{\gamma}{r^2M}\right)^{\frac{1}{2}} \cdot \left(1 - \frac{\gamma}{r^2M}\right)^{r^2M} \cdot (r^2M - \gamma)^{-\gamma} = A_1 \cdot A_2 \cdot A_3,$$

where $\lim_{M \rightarrow \infty} A_1 = 1$ and $\lim_{M \rightarrow \infty} A_2 = e^{-\gamma}$. Similarly, the second fraction decomposes as $B_1 \cdot B_2 \cdot B_3$, where $\lim_{M \rightarrow \infty} B_1 = 1$, $\lim_{M \rightarrow \infty} B_2 = e^{2\gamma}$ and $B_3 = ((r - r^2)M + \gamma)^{2\gamma}$. The third fraction can be decomposed as $C_1 \cdot C_2 \cdot C_3$, where $\lim_{M \rightarrow \infty} C_1 = 1$, $\lim_{M \rightarrow \infty} C_2 = e^{-\gamma}$ and $C_3 = ((1 - r)^2M - \gamma)^{-\gamma}$. Thus, the limit of expression (2) as $M \rightarrow \infty$ reduces to

$$\begin{aligned} \lim_{M \rightarrow \infty} \frac{1}{e^{\gamma(r^2M - \gamma)^{\gamma}}} \cdot e^{2\gamma((r - r^2)M + \gamma)^{2\gamma}} \cdot \frac{1}{e^{\gamma((1 - r)^2M - \gamma)^{\gamma}}} &= \\ \lim_{M \rightarrow \infty} \left(\frac{((r - r^2)M + \gamma)^2}{(r^2M - \gamma)((1 - r)^2M - \gamma)} \right)^{\gamma} &= 1, \end{aligned}$$

which delivers the desired result. \square

Proof of Proposition 5. First, I know that at most $M - \alpha - 1$ people were infected after period \bar{t} . The new limit vector $y^* \in \mathbb{R}^{M - \alpha - 1}$ must be computed using matrix $\bar{C}_{(\alpha+1)}$. If we define $z := (\bar{y}_1^M, \dots, \bar{y}_{M - \alpha - 1}^M)$, it is easy to see that $y^* = \frac{z}{\|z\|}$.³⁹

Second, consider the following scenario. Suppose that, late in Phase III, an infected player happened to believe that exactly two people were infected in each community, and then he played the Nash action for multiple periods while observing only g . In each period, he infected a new person and contagion continued to spread exponentially. If the number of periods in which the player infected people is large enough, Nash reversion would be the best reply. This number of periods only depends on the specific game G and on the population size M and, thus, we denote this number by $\phi^G(M)$. Since the contagion spreads

³⁹One way to see this is to recall that $\bar{y}^M = \lim_{t \rightarrow \infty} \frac{\bar{y}^t \bar{C}_2}{\|\bar{y}^t \bar{C}_2\|}$, $y^* = \lim_{t \rightarrow \infty} \frac{z^t \bar{C}_{(\alpha+1)}}{\|z^t \bar{C}_{(\alpha+1)}\|}$ and note that, for each t and each $k \in \{1, \dots, M - \alpha - 1\}$, $z_k^t = \frac{\bar{y}_k^t}{\|\bar{y}^t\|}$.

exponentially, for fixed G , $\phi^G(M)$ is some logarithmic function of M . Hence, for each $\hat{r} \in (0, 1)$, there is \hat{M} such that for $M > \hat{M}$, $\hat{r}M > \phi^G(M)$. Now, given $r \in (0, 1)$ and $m \in \mathbb{N}$, we can find $\hat{r} \in (r, 1)$ and \underline{M} such that Lemma 5 holds. For the rest of the proof, we work with $M > \max\{\hat{M}, \underline{M}\}$. There are two cases:

$\alpha < \phi^G(M)$: In this case, we can repeat the arguments in the proof of Proposition 4 to show that my beliefs $x^{\bar{t}+1+\alpha}$ first-order stochastically dominate y^* . Since I observed α periods of good behavior after being infected, I know that at most $M - \alpha$ players were infected at the end of period \bar{t} . Now, since $\hat{r}M > \phi^G(M)$, we have that $\lfloor M - \hat{r}M \rfloor < M - \phi^G(M) < M - \alpha$. Then, we can parallel the arguments in Proposition 4 and use Lemma 5 to find an M_G that ensures that the player has strict incentives to play the Nash action if the population size is at least M_G .

$\alpha \geq \phi^G(M)$: In this case, I played the Nash action for α periods. By definition of $\phi^G(M)$, playing Nash is the unique best reply after observing $h^{\bar{t}+1+\alpha}$.

Finally, we can just take $M_G^2 = \max\{M_G, \hat{M}, \underline{M}\}$. □

B.5 Incentives of Patient Players

We formally show here that even an infinitely patient player would have the incentive to punish when prescribed to do so. This is important for the discussion in Section A.4.

Let \bar{m} be the maximum possible gain any player can make by a unilateral deviation from any action profile of the stage-game. Let \underline{l} be the minimum loss that a player suffers in the stage-game when he does not play his best response to his rival's strict Nash action a^* (recall that, since a^* is a strict Nash equilibrium, $\underline{l} > 0$). Suppose that we are in Phase III, and take a player i who knows that the contagion has started. For the analysis in this section, we use the continuation payoffs of the Nash reversion strategy as the benchmark.

For any given continuation strategy of a player i , let $v(M)$ denote player i 's (expected) undiscounted sum of future gains relative to his payoff from the Nash reversion strategy. It is easy to see that $v(M)$ is finite. Player i knows that contagion is spreading exponentially and, hence, his payoff will drop to the Nash payoff $u(a^*)$ in the long run. In fact, although $v(M)$ increases with M , since contagion spreads exponentially fast, $v(M)$ grows at a slower rate than M . Similarly, for any continuation strategy, and for each $r \in (0, 1]$, let $v(r, M)$ denote the (expected) undiscounted sum of future gains player i can make from playing this continuation strategy relative to the payoff from Nash reversion, when he is in Phase III and knows that at least rM people are infected in each community. In the result

below, we show that $v(r, M)$ is uniformly bounded on M .

Lemma 6. *Fix a game $G \in \mathcal{G}$. Let $r > 1/2$. Then, there is \bar{U} such that, for each $\bar{r} > r$ and each $M \in \mathbb{N}$, $v(\bar{r}, M) \leq \bar{U}$.*

Proof. Let $P(r, M)$ denote the probability of the event: “If $\lceil rM \rceil$ people are infected at period t , the number of infected people at period $t + 1$ is, at least, $\lceil rM + \frac{(1-r)}{2}M \rceil$,” that is, $P(r, M)$ is the probability that at least half the uninfected people get infected in the present period, given that rM people are infected already. We claim that $P(\frac{1}{2}, M) \geq \frac{1}{2}$.

Suppose that M is even and $\frac{M}{2}$ people are infected at the start of period t (i.e., $r = \frac{1}{2}$). What is the probability of at least $\frac{M}{4}$ more people getting infected in this period? It is easy to see from the contagion matrix that there is a symmetry in the transition probabilities of different states. The probability of no one getting infected in this period is the same as that of no one remaining uninfected. The probability of one person being infected is the same as that of only one person remaining uninfected. In general, probability of $k < M/4$ players getting infected is the same as that of k players remaining uninfected. This symmetry implies immediately that $P(\frac{1}{2}, M) \geq \frac{1}{2}$.

It is easy to see that for $r > \frac{1}{2}$, $\lceil rM \rceil > \frac{M}{2}$, and further $P(r, M) > P(\frac{1}{2}, M) \geq \frac{1}{2}$. Thus, for each $r > \frac{1}{2}$ and each $M \in \mathbb{N}$, $P(r, M) > \frac{1}{2}$. Also, for any fixed $r > \frac{1}{2}$, $\lim_{M \rightarrow \infty} P(r, M) = 1$. Intuitively, if more than half the population is already infected, then the larger M is, the more unlikely it is that more than half of the uninfected people remain uninfected. Hence, given $r > \frac{1}{2}$, there is $\hat{p} < \frac{1}{2}$ such that, for each $M \in \mathbb{N}$, $P(r, M) > 1 - \hat{p}$.

Given any continuation strategy, we want to compute $v(r, M)$. Note that if a player meets an infected opponent, then not playing the Nash action involves a minimal loss of \underline{l} ; and if he meets an uninfected opponent, his maximal gain relative to playing the Nash action is \bar{n} . So, we first compute the probability of meeting an uninfected player in any future period. Suppose that $\lceil rM \rceil$ people are infected at the start of period t . Let $s_0 = 1 - r$.

Current period t : There are $\lceil rM \rceil$ infected players. So, the probability of meeting an uninfected opponent is $(1 - r) = s_0$.

Period $t + 1$: With probability at least $1 - \hat{p}$, at least half the uninfected players got infected in period t and with probability at most \hat{p} , less than half the uninfected players got infected. So, the probability of meeting an uninfected opponent in this period is less than $(1 - \hat{p})\frac{s_0}{2} + \hat{p}s_0 < \frac{s_0}{2} + \hat{p}s_0$.

Period $t + 2$: Similarly, using the upper bound $\frac{s_0}{2} + \hat{p}s_0$ obtained for period $t + 1$, the probability of meeting an uninfected opponent in period $t + 2$ is less than

$$(1 - \hat{p})\left(\frac{s_0}{4} + \hat{p}\frac{s_0}{2}\right) + \hat{p}\left((1 - \hat{p})\frac{s_0}{2} + \hat{p}s_0\right) \leq \frac{s_0}{4} + 2\hat{p}\frac{s_0}{2} + \hat{p}^2s_0 = s_0\left(\frac{1}{2} + \hat{p}\right)^2.$$

Period $t + \tau$: Probability of meeting an uninfected opponent in period $t + \tau$ is less than $s_0\left(\frac{1}{2} + \hat{p}\right)^\tau$.

Therefore, we have

$$\begin{aligned} v(r, M) &< s_0\bar{m} - (1 - s_0)\underline{l} + s_0\left(\frac{1}{2} + \hat{p}\right)\bar{m} - (1 - s_0\left(\frac{1}{2} + \hat{p}\right))\underline{l} + \dots \\ &= \sum_{\tau=0}^{\infty} s_0\left(\frac{1}{2} + \hat{p}\right)^\tau\bar{m} - (1 - s_0\left(\frac{1}{2} + \hat{p}\right)^\tau)\underline{l} \leq \sum_{\tau=0}^{\infty} s_0\left(\frac{1}{2} + \hat{p}\right)^\tau\bar{m} \\ &= s_0\bar{m} \sum_{\tau=0}^{\infty} \left(\frac{1}{2} + \hat{p}\right)^\tau = \bar{U}. \end{aligned}$$

Convergence of the series follows from the fact that $\frac{1}{2} + \hat{p} < 1$. Clearly, for $\bar{r} > r$, $v(\bar{r}, M) \leq v(r, M)$, and, hence, $v(\bar{r}, M)$ is uniformly bounded, *i.e.*, for each $\bar{r} \geq r$ and each $M \in \mathbb{N}$, $v(\bar{r}, M) \leq \bar{U}$.⁴⁰ \square

Proposition 7. *Fix a game $G \in \mathcal{G}$. Then, there exist $\underline{r} \in (0, 1)$ and $\underline{M} \in \mathbb{N}$ such that, for each $r \geq \underline{r}$ and each $M \geq \underline{M}$, a player who gets infected very late will not slow down the contagion even if he is infinitely patient ($\delta = 1$).*

Proof. Take $r > \frac{1}{2}$. By Lemma 4, if M is big enough, a player who gets infected late in Phase III believes that “with probability at least $1 - \frac{1}{M^2}$, at least rM people in each community are infected.” Consider a continuation strategy where he deviates and does not play the Nash action.

- i) With probability $1 - \frac{1}{M^2}$, at least rM people are infected. So, with probability at least r , he meets an infected player, makes a loss of at least \underline{l} by not playing Nash,

⁴⁰The upper bound obtained here is a very loose one. Notice that contagion proceeds at a very high rate. In the proof of Lemma 4, we showed that if the state of contagion is such that number of infected people is rM , the most likely state in the next period is $(2r - r^2)M$. This implies that the fraction of uninfected people goes down from $(1 - r)$ to $(1 - 2r - r^2) = (1 - r)^2$. More generally, if we consider contagion evolving along this path of “most likely” transitions for t consecutive periods, the number of uninfected people would go down to $(1 - r)^{2^t}$, *i.e.*, the contagion spreads at a highly exponential rate. In particular, this rate of convergence is independent of M .

and does not slow down the contagion. With probability $1 - r$, he gains, at most, \bar{m} in the current period and $v(r, M)$ in the future.

- ii) With probability $\frac{1}{M^2}$, fewer than rM people are infected, and the player's gain is, at most, \bar{m} in the current period and, at most, $v(M)$ in the future.

Hence, by Lemma 6, the gain from not playing the Nash action instead of doing so is bounded above by:

$$\frac{\bar{m} + v(M)}{M^2} + \left(1 - \frac{1}{M^2}\right) (-rl + (1-r)(\bar{m} + v(M, r))) < \frac{M + \bar{m}}{M^2} + \left(1 - \frac{1}{M^2}\right) (-rl + (1-r)(\bar{m} + \bar{U})).$$

The inequality follows from the facts that $v(M)$ is finite and increases slower than the rate of M and that \bar{U} is a uniform bound for $v(r, M)$ for any $r > \frac{1}{2}$. If M is large enough and r is close to 1, the expression is negative. So, there is no incentive to slow down the contagion. \square

B.6 Uncertainty about Timing

In this section we investigate what happens if players are not sure about the the timing of the different phases. We conjecture that a modification of our strategies is robust to the introduction of *small* uncertainty about timing. To provide some intuition for this conjecture, we consider a setting in which there is uncertainty about the timing of the phases. More precisely, the players don't know precisely which equilibrium they are supposed to play, among all that can be defined via pairs (\dot{T}, \ddot{T}) . For the sake of exposition, we restrict attention to the product-choice game and try to sustain a payoff close to the efficient outcome $(1, 1)$.

Given the product-choice game and community size M , we choose a pair \dot{T} and \ddot{T} that induces an equilibrium. At the start of the game, each player receives an independent, noisy but informative signal about the timing of the trust-building phases (values of \dot{T} and \ddot{T}). Each player receives a signal $\omega_i = (\dot{d}_i, \dot{\Delta}_i, \ddot{d}_i, \ddot{\Delta}_i)$, which is interpreted as follows. Player i , on receiving signal ω_i , can bound the values of \dot{T} and \ddot{T} with two intervals; *i.e.*, she knows that $\dot{T} \in [\dot{d}_i, \dot{d}_i + \dot{\Delta}_i]$ and $\ddot{T} \in [\ddot{d}_i, \ddot{d}_i + \ddot{\Delta}_i]$. The signal generation process is described below. The idea is that players are aware that there are two trust-building phases followed by the target payoff phase, but are uncertain about the precise timing of the phases. Moreover, signals are informative in that the two intervals are non-overlapping and larger intervals (imprecise estimates) are less likely than smaller ones.

- i) $\dot{\Delta}_i$ is drawn from a Poisson distribution with parameter $\dot{\gamma}$, and then \dot{d}_i is drawn from the discrete uniform distribution over $[\dot{T} - \dot{\Delta}_i, \dot{T} - \dot{\Delta}_i + 1, \dots, \dot{T}]$. If either 1 or \dot{T} lie in the resulting interval $[\dot{d}_i, \dot{d}_i + \dot{\Delta}_i]$, then $\dot{\Delta}_i$ and \dot{d}_i are drawn again.
- ii) After $\dot{\Delta}_i$ and \dot{d}_i are drawn as above, $\ddot{\Delta}_i$ is drawn from a Poisson distribution with parameter $\ddot{\gamma}$. Finally, \ddot{d}_i is drawn from the discrete uniform distribution over $[\ddot{T} - \dot{\Delta}_i, \ddot{T} - \dot{\Delta}_i + 1, \dots, \ddot{T}]$. If the resulting interval, $[\ddot{d}_i, \ddot{d}_i + \ddot{\Delta}_i]$ overlaps with the first interval $[\dot{d}_i, \dot{d}_i + \dot{\Delta}_i]$, (*i.e.*, $\dot{d}_i + \dot{\Delta}_i \in [\ddot{d}_i, \ddot{d}_i + \ddot{\Delta}_i]$), then \ddot{d}_i is redrawn.

In this setting, players are always uncertain about the start of the trust-building phases and precise coordination is impossible. However, we conjecture that with a modification to our strategies, sufficiently patient players will be able to attain payoffs arbitrarily close to $(1, 1)$, if the uncertainty about timing is very small.

Equilibrium play: Phase I: Consider any player i with signal $\omega_i = (\dot{d}_i, \dot{\Delta}_i, \ddot{d}_i, \ddot{\Delta}_i)$. During the first $\dot{d}_i + \dot{\Delta}_i$ periods, he plays the cooperative action (Q_H or B_H). **Phase II:** During the next $\ddot{d}_i - (\dot{d}_i + \dot{\Delta}_i)$ periods, he plays as if he were in Phase II, *i.e.*, a seller plays Q_L and a buyer B_H . **Phase III:** For the rest of the game (*i.e.*, from period \ddot{d}_i on), he plays the efficient action (Q_H or B_H).

Off-Equilibrium play: A player can be in one of two moods: *uninfected* or *infected*, with the latter mood being irreversible. We define the moods a little differently. At the beginning of the game, all players are uninfected. Any action (observed or played) that is not consistent with play that can arise on-path, given the signal structure, is called a deviation. We classify deviations into two types. Deviations that definitely entail a short-run loss for the deviating player are called *non-triggering* deviations (e.g., a buyer deviating in the first period of the game). Any other deviation is called a *triggering* deviation (*i.e.*, these are deviations that with positive probability give the deviating player a short-run gain). A player who is aware of a triggering deviation is said to be infected. Below, we specify off-path behavior. We do not completely specify play after all histories, but the description below will suffice to provide the intuition behind the conjecture. An uninfected player continues to play as if on-path. An infected player acts as follows.

- Deviations observed before $\dot{d}_i + \dot{\Delta}_i$: A buyer i who gets infected before period \dot{d}_i switches to her Nash action forever at some period between \dot{d}_i and $\dot{d}_i + \dot{\Delta}_i$ when

she believes that enough buyers are infected and have also switched. Note that buyers cannot get infected between \dot{d}_i and $\dot{d}_i + \dot{\Delta}_i$, since any action observed during this period is consistent with equilibrium play (*i.e.*, a seller j playing Q_L at time $t \in [\dot{d}_i, \dot{d}_i + \dot{\Delta}_i]$ may have received a signal such that $\dot{d}_j + \dot{\Delta}_j = t$).

A seller i who faces B_L before period \dot{d}_i , ignores it (this is a non-triggering deviation, as the buyer must still be in Phase I, which means that the deviation entails a short-term loss for her). If a seller observes B_L between periods \dot{d}_i and $\dot{d}_i + \dot{\Delta}_i$, he will switch to Nash immediately.

- Deviations observed between $\dot{d}_i + \dot{\Delta}_i + 1$ and \ddot{d}_i : A player who gets infected in the interval $[\dot{d}_i + \dot{\Delta}_i + 1, \ddot{d}_i]$ will switch to Nash forever from period \ddot{d}_i . Note that buyers who observe Q_H ignore such deviations as they are non-triggering.
- Deviations observed after \ddot{d}_i : A player who gets infected after \ddot{d}_i switches to the Nash action immediately and forever.

We argue below why these strategies can constitute an equilibrium by analyzing some important histories. **Incentives of players on-path:** If triggering deviations are definitely detected and punished by Nash reversion, then, for sufficiently patient players, the short-run gain from a deviation will be less than the long-term loss in payoff from starting the contagion. So, we need to check that all deviations are detected (though, possibly with probability < 1), and that the resultant punishment is enough to deter the deviation.

- Seller i deviates (plays Q_L) at $t = 1$: With probability 1, his opponent will detect the deviation, and ultimately his payoffs will drop to a very low level. A sufficiently patient player will, therefore, not deviate.
- Seller i deviates at $2 \leq t < \dot{d}_i + \dot{\Delta}_i$: With positive probability, his opponent j has $\dot{d}_j > t$, and will detect the deviation and punish him. But, because of the uncertainty about the values of \dot{T} and \ddot{T} , with positive probability, the deviation goes undetected and unpunished. The probability of detection depends on the time of the deviation (detection is more likely earlier than later, because early on, most players are outside their first interval). So, the deviation gives the seller a small current gain with probability 1, but a large future loss (from punishment) with probability less than 1. If the uncertainty about \dot{T} and \ddot{T} is small enough (*i.e.*, signals are very precise), then the probability of detection (and future loss) will be high. For a sufficiently patient player, the current gain will then be outweighed by the expected future loss.

- Seller i deviates (plays Q_L) at $t \geq \dot{d}_i$: With positive probability, his opponent j has signal $\ddot{d}_j = \dot{d}_i$, and will detect the deviation.
- All deviations by buyers (playing B_L) are detected, since B_L is never consistent with equilibrium play. If a buyer plays a triggering deviation B_L , she knows that with probability 1, her opponent will start punishing immediately. The buyer's incentives in this case are exactly as in the setting without uncertainty. For appropriately chosen \dot{T} and \ddot{T} , buyers will not deviate on-path.

Optimality of Nash reversion off-path: Now, because players are uncertain about the true values of \dot{T} and \ddot{T} , there are periods when they cannot distinguish between equilibrium play and deviations. We need to consider histories where a player can observe a triggering deviation, and check that it is optimal for him to start punishments. As before, we assume that players, on observing a deviation, believe that some seller deviated in the first period.

First, consider incentives of a seller i . We argue that a seller who deviates at $t = 1$ will find it optimal to continue deviating. Further, a seller who gets infected by a triggering deviation at any other period will find it optimal to revert immediately to the Nash action.

- Suppose that seller i deviates at $t = 1$, and plays Q_L . He knows that his opponent will switch to the Nash action at most, at the end of her first interval (close to the true \dot{T} with high probability), and the contagion will spread exponentially from some period close to the true $\dot{T} + \ddot{T}$. Thus, if seller i is sufficiently patient, his continuation payoff will drop to a very low level after $\dot{T} + \ddot{T}$, regardless of his play in his Phase I (until period $\dot{d}_i + \dot{\Delta}_i$). Therefore, for a given M , if \dot{T} is large enough (and so $\dot{d}_i + \dot{\Delta}_i$ is large), the optimal continuation strategy for seller i will be to continue playing Q_L .
- Seller i observes a triggering deviation of B_L : If a seller observes a triggering deviation of B_L by a buyer (in Phase II), he thinks that the first deviation occurred at period 1, and by now all buyers are infected. Since, his play will have a negligible effect on the contagion process, it is optimal to play Q_L .

Now, consider the incentives of a buyer.

- Buyer i observes Q_L at $1 \leq t < \dot{d}_i$: This must be a triggering deviation. A seller j should switch to Q_L only at the end of his first interval ($\dot{d}_j + \dot{\Delta}_j$), and this cannot be the case because, then, the true \dot{T} does not lie in player i 's first interval. On observing this triggering deviation, the buyer believes that the first deviation occurred at $t = 1$

and the contagion has been spreading since then. Consequently, she will switch to her Nash action forever at some period between \dot{d}_i and $\dot{d}_i + \dot{\Delta}_i$ when she begins believing that enough other buyers are infected and have switched, as well (It is easily seen that, at worst, buyer i will switch at period $\dot{d}_i + \dot{\Delta}_i$).

- Buyer i observes Q_L at $t \geq \ddot{d}_i + \ddot{\Delta}_i$. Since i is at the end of her second interval, she knows that every rival has started his second interval, and should be playing Q_H . So, this is a triggering deviation. She believes that the first deviation occurred at $t = 1$, and so most players must be infected by now. This makes Nash reversion optimal.

Note that in any other period, buyers cannot distinguish a deviation from equilibrium play.

- i) Any action observed by buyer i in her first interval (*i.e.*, for t such that $\dot{d}_i \leq t < \dot{d}_i + \dot{\Delta}_i$) is consistent with equilibrium play. A seller j playing Q_H could have got signal $\dot{d}_j > t$, and a seller playing Q_L could have got signal $\dot{d}_j + \dot{\Delta}_j \leq t$.
- ii) Any action observed by buyer i between her two intervals (*i.e.*, at t such that $\dot{d}_i + \dot{\Delta}_i \leq t < \ddot{d}_i$) is consistent with equilibrium play. Q_L is consistent with a seller j who got $\dot{d}_j + \dot{\Delta}_j \leq t$, and Q_H is consistent with a seller with signal such that $t < \dot{d}_j + \dot{\Delta}_j$.
- iii) Any action observed by buyer i within her second interval (*i.e.*, at t such that $\ddot{d}_i \leq t < \ddot{d}_i + \ddot{\Delta}_i$) is consistent with equilibrium play. Q_L is consistent with a seller j who got $\ddot{d}_j > t$ (say $\ddot{d}_j = \ddot{d}_i + \ddot{\Delta}_i$), and Q_H is consistent with a seller with signal such that $\ddot{d}_j < t$ (say j got the same signal as buyer i).

B.7 Interchangeable Populations

We assumed that each player belongs either to Community 1 or to Community 2. Alternatively, we could have assumed that there is one population whose members are matched in pairs in every period and, in each match, the roles of players are randomly assigned. At the start of every period, each player has an equal probability of playing as player 1 or 2 in the stage-game. A first implication of this setting is that a negative result like Proposition 1 may no longer be true.⁴¹ We believe that the trust-building ideas that underlie this paper can be adapted to this new setting.

⁴¹A buyer infected in period 1 might become a seller in period 2, and he might indeed have the right incentives to punish.

Consider the repeated product-choice game when roles are randomly assigned in each period. We conjecture that the following version of our strategies can be used to get as close as desired to payoff $(1, 1)$. There are two phases. Phase I is the trust-building phase: sellers play Q_L and buyers play B_H ; the important features of this profile being that i) only buyers have an incentive to deviate, and ii) sellers are playing a Nash action. Phase II is the target payoff phase and (Q_H, B_H) is played. Deviations are punished through Nash reversion with no delay in the punishment. Thus, contagion also takes place in Phase I; whenever an “infected” player is in the role of a buyer, he will play B_L and spread the contagion, so we do not have a single player infecting people in this phase. This implies that we do not need a second trust-building phase, since its primary goal was to give the infected buyers the right incentives to “tell” the sellers that there had been a deviation.

The arguments would be very similar to those in the setting with independent populations. After getting infected, a player would believe that a buyer deviated in period one and that punishments have been going on ever since. The key observation is that an infected player early in Phase I knows that contagion will spread regardless of his action, and so he wants to get all the short-run gains he can when he is a buyer (similar to the incentives of the deviant seller in Phase I in our construction). Proving formally that players have the right incentives after all histories is hard, and we cannot rely on the analysis of the independent populations. Interchangeable roles for players have two main consequences for the analysis. First, the contagion process is not the same, and a slightly different mathematical object is needed to model it. Second, the set of histories a player observes would depend on the roles he played in the past periods, so it is harder to characterize all possible histories. We think that this exercise would not yield new insights, and so leave it as a conjecture.

B.8 Can we get a (Nash Threats) Folk Theorem?

Our strategies do not suffice to get a folk theorem for all games in \mathcal{G} . For a game $G \in \mathcal{G}$ with strict Nash equilibrium a^* , the set F_{a^*} does not include action profiles where only one player is playing the Nash action a_i^* . For instance, in the product-choice game, our construction cannot achieve payoffs close to $(1 + g, -l)$ or $(-l, 1 - c)$. However, we believe that the trust-building idea is powerful enough to take us farther. We conjecture we can obtain a Nash threats folk theorem for two-player games by modifying our strategies with the addition of further trust-building phases. We do not prove a folk theorem here, but hope that the informal argument below will illustrate how this might be done. To fix ideas,

we restrict attention to the product-choice game, although the extension to general games may entail additional difficulties.

Consider a feasible and individually rational target payoff that can be achieved by playing short sequences of (Q_H, B_H) (10 percent of the time) alternating with longer sequences of (Q_H, B_L) (90 percent of the time). It is not possible to sustain this payoff in Phase III with our strategies. To see why not, consider a long time window in Phase III where the prescribed action profile is (Q_H, B_L) . Suppose that a buyer faces Q_L for the first time in a period of this phase followed by many periods of Q_H . Notice that since the action for a buyer is B_L in this time window, she cannot infect any sellers herself. Then, with more and more observations of Q_H , she will ultimately be convinced that few people are infected. Thus, it may not be optimal to keep playing Nash any more. This is in sharp contrast with the original situation, where the target action was (Q_H, B_H) . In that case, a player who gets infected starts infecting players himself and so, after, at most, $M - 1$ periods of infecting opponents, he is convinced that everyone is infected.

What modification to our strategies might enable us to attain these payoffs? We can use additional trust-building phases. Consider a target payoff phase that involves alternating sequences of (Q_H, B_L) for T_1 periods and (Q_H, B_H) for $T_2 = \frac{1}{9}T_1$ periods. In the modified equilibrium strategies, in Phase III, the windows of (Q_H, B_L) and (Q_H, B_H) will be separated by trust-building phases. To illustrate, we start the game as before, with two phases: \dot{T} periods of (Q_H, B_H) and \ddot{T} periods of (Q_L, B_H) . In Phase III, players play the action profile (Q_H, B_L) for T_1 periods, followed by a new trust-building phase of T' periods during which (Q_L, B_H) is played. Then, players switch to playing the sequence of (Q_H, B_H) for T_2 periods. The new phase is chosen to be short enough (*i.e.*, $T' \ll T_1$) to have no significant payoff consequences. But, it is chosen long enough so that a player who is infected during the T_1 period window, but thinks that very few people are infected, will still want to revert to Nash punishments to make short-term gains during the new phase.⁴² We conjecture that adding appropriate trust-building phases in the target payoff phase in this way can guarantee that players have the incentive to revert to Nash punishments off-path for any beliefs they may have about the number of people infected.

⁴²For example, think of a buyer who observes a triggering action for the first time in Phase III (while playing (Q_H, B_L)) and then observes only good behavior for a long time while continuing to play (Q_H, B_L) . Even if this buyer is convinced that very few people are infected, she knows that the contagion has begun, and ultimately her continuation payoff will become very low. So, if there is a long enough phase of playing (Q_L, B_H) ahead, she will choose to revert to Nash because this is the myopic best response, and would give her at least some short-term gains.

B.9 A Problematic Game outside \mathcal{G}

Consider the two-player game in Figure 5. This is a game with strictly aligned interests. For

	L	C	R
T	$-5, -5$	$-1, 8$	$5, 5$
M	$-5, -5$	$-2, -2$	$8, -1$
B	$-3, -3$	$-5, -5$	$-5, -5$

Figure 5: A game outside \mathcal{G} .

each (pure) action profile, either it is a Nash equilibrium or both players want to deviate. The difference with respect to other strictly aligned interests games, such as the battle of the sexes or the chicken game, is that there is a Pareto efficient payoff, $(5, 5)$, that cannot be achieved as the convex combination of Nash payoffs. Further, since it Pareto dominates the pure Nash given by (B, L) , it might be possible to achieve it using Nash reversion. Note that, given a strictly aligned interests game and an action profile, if a player plays her best reply against her opponent's action, the resulting profile is a Nash equilibrium.

What is special about games like this that make it harder to get cooperation? Suppose that we want to get close to payoff $(5, 5)$ in equilibrium in the repeated random matching game. Both players have a strict incentive to deviate from (T, R) , the action profile that delivers $(5, 5)$. For this game, we cannot rely on the construction we used to prove Theorem 1, since there is no one-sided incentive profile to use in Phase I.

The approach of starting play with an initial phase of playing Nash action profiles does not work well. This suggests that play in the initial phase of the game should consist of action profiles where both players have an incentive to deviate. Suppose that we start the game with a phase in which we aim to achieve target payoff $(5, 5)$, with the threat that any deviation will, at some point, be punished by Nash reversion to $(-3, -3)$. Suppose that a player deviates in period 1. Then, the opponent knows that no one else is infected in her community and that Nash reversion will eventually occur. Hence, both infected players will try to make short-run gains by moving to the profile that gives them 8. As more players become infected, more people are playing M and C and the payoff will get closer to $(-2, -2)$. In this situation, it is not clear how the best-replying dynamics will evolve. Further, what is even more complicated, is that it is very hard to provide the players with the right incentives to move to $(-3, -3)$. Note that, as long as no player plays B or L , no one ever gets something below -2 , while B and L lead to, at most, -3 . That is, a player

will not switch to B unless she thinks that a high number of players in the other community are already playing L but, who will be the first one to switch? Can we ensure somehow that players coordinately switch to $(-3, -3)$ in this setting?