

Supplement to “Enforcing social norms: Trust-building and community enforcement”

(*Theoretical Economics*, Vol. 14, No. 4, November 2019, 1387–1433)

JOYEE DEB

Department of Economics, Yale University

JULIO GONZÁLEZ-DÍAZ

Department of Economics, University of Santiago de Compostela

APPENDIX B

B.1 *Updating of beliefs conditional on observed histories*

We validate the approach to computing beliefs discussed in Section 4.2. Suppose that player i observes history $h^{\bar{t}+1} = g \dots gbg$ in Phase III. We want to compute her beliefs at the end of period $\bar{t} + 1$ conditional on $h^{\bar{t}+1}$, namely $x^{\bar{t}+1}$. We first compute a set of intermediate beliefs x^t for $t < \bar{t} + 1$. For any period $t < \bar{t}$, we compute x^{t+1} from x^t by conditioning on G^{t+1} and $\mathcal{U}^{t+1} \leq M - 2$. We do not use the information that “I was healthy at the end of each period t^* with $t + 1 < t^* < \bar{t}$.” This information is added later, period by period, i.e., only at period t do we add the information that “I was healthy at the end of period t .” We show that this method is equivalent to 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 $g \dots gbg$. ^{α} . Let b^t (g^t) denote the event “I faced b (g) in period t .” Moreover, the following statements hold.

- The term $U_{i,k}^t$ denotes the event $i \leq \mathcal{U}^t \leq k$, i.e., the number of unhealthy sellers at the end of period t is at least i and at most k .
- We have $E_\alpha^t := U_{1, M-\alpha-1}^t \cap G^t$.
- We have $E_\alpha^{t+1} := E_\alpha^t \cap U_{2, M-\alpha}^{t+1} \cap b^{t+1}$.
- For each $\beta \in \{1, \dots, \alpha - 1\}$, $E_\alpha^{t+1+\beta} := E_\alpha^{t+\beta} \cap U_{\beta+2, M-\alpha+\beta}^{t+1+\beta} \cap g^{t+1+\beta}$.
- We have $E_\alpha^{t+1+\alpha} := E_\alpha^{t+\alpha} \cap U_{\alpha+2, M}^{t+1+\alpha} \cap g^{t+1+\alpha} = h^{t+1+\alpha}$.

Let H^t be a complete history of the contagion process up to period t . Let \mathcal{H}^t be the set of all H^t histories. Let $\mathcal{H}_k^t := \{H^t \in \mathcal{H}^t : \mathcal{U}^t = k\}$. We say $H^{t+1} \Rightarrow h^{t+1}$ if, under H^{t+1} , I observed h^{t+1} . Given $\beta \in \{0, \dots, \alpha\}$, let $P(i \xrightarrow{t+1+\beta} k) := P(\mathcal{U}^{t+1+\beta} = k |_{E_\alpha^{t+1+\beta} \cap \mathcal{U}^{t+\beta} = i})$. Since

Joyee Deb: joyee.deb@yale.edu

Julio González-Díaz: julio.gonzalez@usc.es

$E_\alpha^{t+1+\alpha} = h^{t+1+\alpha}$, the probabilities of interest are $P(\mathcal{U}^{t+1+\alpha} = k | E_\alpha^{t+1+\alpha})$. We claim that these probabilities can be obtained by starting with the probabilities after t conditional on E_α^t and then letting the contagion elapse one more period at a time, conditioning on the information “in the current period I observed g and infected one more person.” Formally, we want to show that, for each $\beta \in \{0, \dots, \alpha\}$,

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

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 $H^{t+1+\beta}, \tilde{H}^{t+1+\beta} \in F_k^{1+\beta} \cap F_i^\beta$, $P(H^{t+1+\beta} |_{H^{t+1+\beta, \beta}}) = P(\tilde{H}^{t+1+\beta} |_{\tilde{H}^{t+1+\beta, \beta}})$. 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 $h^{t+1+\alpha} = E_\alpha^{t+1+\beta}$. Clearly, this number is independent of the history that led to i people being unhealthy. Then we have $P(i \xrightarrow{t+1+\beta} k) = P(F_i^\beta \xrightarrow{t+1+\beta} F_k^{1+\beta}) |i \xrightarrow{t+1+\beta} k|$. Therefore,

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

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

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

B.2 Incentives after histories with multiple deviations

We now discuss different types of histories that can arise when multiple deviations occur.

- H1.** Consider the situation in which a rogue player, after his initial deviation, observes a probability 0 history. His behavior has not been specified. We say only that he will best respond given his beliefs. Analogously to point (iii) in the statement of Lemma 2, this rogue player will assign probability 1 to these deviations being errors by infected players. In particular, a rogue seller who deviated in period 1 will not believe that contagion is proceeding slower than if he had not observed these errors.
- H2.** Consider histories in which a seller deviates in period 1 and then he deviates again during Phase I. The behavior of this rogue seller has not been specified completely. However, we show below that we can still check incentives.
- Consider a history of length $\bar{t} < T^1$ in which a seller deviated in period 1 and in all subsequent periods played an action other than the best response or the on-path action. The best response of the seller at this history would be to play his most profitable deviation until the end of Phase I. This is his best response after his first deviation in period 1. Since any off-path action of a seller in Phase I is a triggering action, the effect of these additional deviations on contagion will be the same as if he had played his best response. An

exposed buyer who observes this behavior will think that she is just facing a seller who deviated in period 1 and is continuing to deviate. Thus, this rogue seller's best response from that point onward will remain the same as if he had been best responding throughout. Also, all the exposed buyers would switch to the Nash action at the end of Phase I.

- Consider a history of length $\bar{t} < T^1$ in which a seller deviated in period 1 and in some later periods played the on-path action. Since on-path actions are not triggering, the above argument can no longer be used to characterize the seller's best response. Yet, any exposed buyer observing an on-path action will think that she is facing a healthy seller while the rogue seller is continuing to infect. Since no one attaches positive probability to such behavior by the rogue seller, not specifying the rogue seller's behavior at such histories is not a problem for analyzing other player's incentives.

H3. Suppose I am a healthy player who observes a triggering action and then deviates from the prescribed off-path action. The strategies prescribe that in subsequent play I ignore my own deviation. To see why this is optimal, we briefly discuss the most problematic case: a history in which I have been infected at a period $\bar{t} + 1$ late in Phase III and observed a history $h^{\bar{t}}$ of the form $h^{t+\alpha} = g \dots gbg \cdot^\alpha \cdot g$. Further, suppose that instead of playing Nash, I have played my on-path action after being infected.

The situation is similar to that covered by Proposition 5, but with the difference that after getting infected, I am not spreading the contagion while observing good behavior. How will my beliefs evolve now? We argue below why, regardless of the value of α , I will still believe that contagion is sufficiently spread for me to have the incentive to play Nash. The argument is very similar to that of Case 1 in the proof of Proposition 5.

History $h^{\bar{t}} = g \dots gbg$. After this history, the argument is completely analogous to Case 1 in the proof of Proposition 5. In that proof, when computing the intermediate beliefs at the end of period \bar{t} , it was argued that they first-order stochastically dominate $\tilde{x}^{\bar{t}}$, the beliefs obtained when conditioning on the following information: (i) I observed g and (ii) at most $M - 2$ people are unhealthy after \bar{t} . In particular, we did not use the information that I had infected an opponent in period \bar{t} , which is the only difference between the history at hand and the histories studied in Case 1 in the proof of Proposition 5. Thus, to get the desired incentives, we can rely again on the fact that $\tilde{x}^{\bar{t}}$ is close to $\bar{y}_{B^1}^M$, the limit of the Markov process with transition matrix $\bar{Q}_{2\downarrow}$.

History $h^{\bar{t}} = h^{t+\alpha} = g \dots gbg \cdot^\alpha \cdot g$. We start with intermediate beliefs x^t . Regardless of the value of α , since I am not spreading contagion (I may be meeting the same healthy player in every period since I got infected), I will still think that at most $M - 2$ people were unhealthy at any period $\tau \leq t$. The transition matrix is $\bar{Q}_{2\downarrow}$, and x^t will be close to $\bar{y}_{B^1}^M$. To compute subsequent intermediate beliefs $x^{t+1}, x^{t+2}, \dots, x^{t+\alpha}$, since I know that at least two people in each community were

unhealthy after \bar{t} , I have to use matrix $\bar{Q}_{[1,2]}$, which shifts the beliefs toward more people being unhealthy (relative to the process given by \bar{Q}_2). Therefore, the ensuing process will move from x^t to a limit that first-order stochastically dominates $\bar{y}_{B^1}^M$ in terms of more people being unhealthy, which ensures that I have the incentive to play Nash.

Finally, to study the beliefs after histories in which, after being infected or exposed, I alternate on-path play with the Nash action and I face both g and b , we would have to combine the above arguments with those in Cases 2 and 3 of the proof of Proposition 5.

B.2.1 Pathological histories Finally, we discuss a class of histories that we call *pathological*. They involve multiple nested off-path deviations combined with a sequence of very low probability match realizations or multiple independent deviations. Behavior has not been described at these histories. They have virtually no effect on incentives: We discuss them for completeness.

PH1. For pedagogical reasons, we start with an extreme example, the *special history*.

Phase I. A seller deviates in period 1 and then meets the same buyer in all periods of Phase I. We call these two players the *special seller* and the *special buyer*, respectively. There is no other deviation during Phase I.

Phase II. In each and every period of Phase II, the special seller further deviates by playing an action that is not the Nash action while being again matched with the special buyer in every period.

Checking incentives after this history is specially difficult. The main role of Phase II is to account for histories in which Phase I proceeds as in this special history. After such histories, when Phase II starts, only one buyer and one seller are unhealthy, and only the buyer knows it. The special seller believes that, with very high probability, every buyer is unhealthy. Since both unhealthy and healthy sellers play Nash during Phase II, the special buyer, while playing Nash in Phase II, will think that, with very high probability, she is infecting all sellers (even if she is meeting the special seller in every period). In the special history, however, the special seller is playing something different from the Nash action. Lemma 2 implies that this erroneous behavior should be attributed to infected players. However, the special buyer knows that there is no infected seller. Since deviations by rogue players are more likely than deviations by healthy players (see Section 4.1), the special buyer will know that she is meeting the special seller (and not spreading the contagion).

For most of Phase II, the special buyer will play Nash and keep making short-run profits (even though, most likely, this will spread the contagion). However, once the end of Phase II approaches and she knows that no seller except the special seller is unhealthy (because she always met the special seller), she might start thinking about playing differently given that contagion is not widely spread. Now, as soon as she plays something that no other buyer (infected or healthy) would play, the special seller will realize that

this pathological history has been realized (note that only the special seller has deviated from the strategy profile so, for him, this history has positive probability given his behavior).

This is a history at which the behavior of two players is not specified and both of them know that it has been realized. But this is not a problem for the following reasons:

- (i) Since this special history is so unlikely, no seller will deviate in period 1, hoping for this extremely unlikely history to be realized. Further, even if he has deviated in period 1, he would not be deviating throughout Phase II, hoping to have met the same buyer throughout Phase I and to be meeting her in each and every period of Phase II.
- (ii) It does not affect the incentives of the special buyer at the start of Phase II, since the strategy prescribes that the rogue seller plays Nash and so she attaches probability 0 to the special history being realized.
- (iii) Lemma 2 ensures that no other player—buyer or seller—will ever assign positive probability to the special history being realized. They will always explain erroneous behavior with deviations by infected players.

The above arguments apply not only to the special history, but also to similar histories that involve a special buyer who observes triggering actions in all periods of Phase I and non-Nash actions in most periods of Phase II. An easier argument applies to similar histories in which, during Phase I, a rogue seller observes only off-path behavior. Since deviations by healthy and exposed buyers are equally likely, his beliefs about contagion are unaltered.

Histories at which behavior is left unspecified for some player can be problematic for the analysis of incentives if other players become aware of these histories. More precisely, underspecification is not problematic if the following property holds.

PROPERTY B. For each pair of players i and j , player j will never assign positive probability to any history at which player i 's behavior is unspecified.

We note that Property B does not hold after the special history. But Lemma 2 ensures that no player other than the special seller and the special buyer will assign positive probability to it.

PH2. Consider histories that involve independent deviations by multiple players. Because behavior has not been specified, if these players become aware of the existence of one another, we might violate Property B. We need to consider the following cases.

- (i) Suppose that seller i becomes rogue in period 1 and player j becomes rogue at a later period. Seller i can never become aware of j 's deviation. In addition, even if player j happens to realize that there is another healthy player who played a triggering action, he will attribute it to a deviation by a seller in period 1. Since the continuation play of such a rogue seller is specified, this history is consistent with Property B.

- (ii) Suppose that two players i and j became rogue (independently) after period 1. If either of them, say i , becomes aware of the existence of another rogue player, he will attribute it to a seller having deviated in period 1. Since continuation play for such a rogue seller is specified, there is no problem in computing i 's incentives.

Note that Lemma 2 ensures that no infected player will ever assign positive probability to histories with multiple rogue players.

B.2.2 Detailed outline of off-path histories and specification of behavior Below we provide a list of off-path histories and discuss how we address the potential issues from not specifying behavior.

Off-path histories for a buyer i

- (i) *Buyer i became rogue by playing the first triggering action of the game.* By definition of a triggering action, a buyer i can become rogue by playing the first triggering action of the game only in Phase II or Phase III. The behavior of buyer i is not specified explicitly at these histories. Equilibrium strategies prescribe that buyer i best responds. However, at these histories, Property B holds: Lemma 1 ensures that no player other than i will ever assign positive probability to such a history being realized.
- (ii) *Buyer i became rogue by playing a triggering action that was not the first triggering action of the game.* By definition, a buyer i can become rogue by playing a triggering action only in Phase II or Phase III. The behavior of buyer i is not specified at these histories. If this was not the first triggering action of the game, then such histories must involve two or more healthy players becoming rogue independently. These histories are *pathological* and have been discussed in Appendix B.2.1.
- (iii) *Buyer i got infected or exposed by facing a triggering action.* The behavior of buyer i is fully specified at these histories. Buyer i ignores the deviation while she is in the exposed mood and switches to the Nash action when she is in the infected mood. However, there are again some *pathological* histories, which were discussed in Appendix B.2.1, where special care is needed to check incentives. This includes, for instance, histories in which, during Phase I, buyer i observes many instances of a seller playing actions that are neither the on-path action nor the prescribed off-path action.

Off-path histories for a seller i

- (iv) *Seller i became rogue by playing the first triggering action of the game.*
 - (a) *Histories in which seller i became rogue by playing the first triggering action of the game in a period $t \neq 1$.* The behavior of seller i is not specified explicitly

at these histories, but the situation is analogous to (i) above, i.e., Property B holds: Lemma 1 ensures that no player other than i will ever assign positive probability to such a history being realized.

(b) *Histories in which seller i became rogue by playing the first triggering action of the game in period 1:*

- Suppose that seller i does not further deviate during Phase I: Behavior of seller i has been specified at these histories. With the exception of the special histories discussed in Appendix B.2.1, no matter what he observes or does, his best response from Phase II onward will be to play the Nash action.
- Suppose that seller i deviates further during Phase I, but does not play the on-path action in any period of Phase I: Behavior of seller i has been specified at these histories. No matter what he observes or does, his best response from Phase II onward will be to play the Nash action. These histories have been discussed in Appendix B.2.1.
- Suppose that seller i deviates further during Phase I, and plays the on-path action at least once in Phase I: The behavior of seller i is not specified. We just prescribe that seller i best responds. Notice that after playing the on-path action for many periods during Phase I, it may no longer be optimal for the seller to keep playing his most profitable deviation throughout Phase I. However, Property B holds at these histories, since any buyer who observes an on-path action in Phase I will believe that she is facing a healthy seller.

(v) *Seller i became rogue by playing a triggering action that was not the first triggering action of the game.* The behavior of seller i is not specified at these histories. Furthermore, at some of these histories, some care is needed to verify that Property B holds. These histories are *pathological* and have been discussed in Appendix B.2.1.

(vi) *Histories in which seller i got infected by facing a triggering action.* The behavior of seller i is fully specified at these histories. He switches to the Nash action forever from the next period.

B.3 Can we get a (Nash threats) folk theorem?

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^* . In the product-choice game, our construction cannot achieve payoffs close to $(1 + g, -l)$ or $(-l, 1 - c)$. However, we conjecture that we can obtain a Nash threats folk theorem for two-player games by modifying our strategies by adding trust-building phases. We hope that the informal argument below illustrates how this might be done in the product-choice game.

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 different from when the target action is (Q_H, B_H) . In that case, a player who gets infected starts infecting players himself and so, after at most $M - 1$ periods, he is convinced that everyone is infected.

Consider a modification: Suppose that the target payoff phase 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. Now, in Phase III, the windows of (Q_H, B_L) and (Q_H, B_H) will be separated by trust-building phases. We start the game as before: T^I periods of (Q_H, B_H) and T^{II} periods of (Q_L, B_H) . In Phase III, players play (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 (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 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 play Nash to make short-term gains during the new phase.* We conjecture that adding such appropriate trust-building phases in the target payoff phase can help obtain a folk theorem.

B.4 A game outside \mathcal{G}

Consider the two-player game in Figure 5. This is a game with strictly aligned interests. Each (pure) action profile is either a Nash equilibrium or both players want to deviate. The difference with other strictly aligned interests games, such as the battle of the sexes, is that there is a Pareto efficient payoff, $(5, 5)$, that cannot be achieved as the convex combination of Nash payoffs. Furthermore, since it Pareto dominates the pure Nash

	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} .

*For example, think of a buyer who observes a triggering action for the first time in Phase III 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 drop. So if there is a long enough phase of playing (Q_L, B_H) ahead, she will play Nash because this is the myopic best response and would give her at least some short-term gains.

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. Suppose that we want to achieve an equilibrium payoff close to $(5, 5)$. Our approach does not work well because there is no one-sided incentive profile to use in Phase I. (Both players have an incentive to deviate from (T, R) .)

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 $(-2, -2)$. Now it is not clear how the dynamics will evolve. Furthermore, it is hard to provide players with the 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 . So a player will not switch to B unless she thinks that many players in the other community are already playing L , but it is not clear who would switch first.

Co-editor George J. Mailath handled this manuscript.

Manuscript received 15 January, 2016; final version accepted 3 April, 2019; available online 15 April, 2019.