

Cooperation and punishment in an adversarial game: How defectors pave the way to a peaceful society

M. B. Short

Department of Mathematics, University of California, Los Angeles, California 90095, USA

P. J. Brantingham

Department of Anthropology, University of California, Los Angeles, California 90095, USA

M. R. D'Orsogna*

Department of Mathematics, California State University at Northridge, Los Angeles, California 91330, USA

(Received 14 September 2010; published 16 December 2010)

The evolution of human cooperation has been the subject of much research, especially within the framework of evolutionary public goods games, where several mechanisms have been proposed to account for persistent cooperation. Yet, in addressing this issue, little attention has been given to games of a more adversarial nature, in which defecting players, rather than simply free riding, actively seek to harm others. Here, we develop an adversarial evolutionary game using the specific example of criminal activity, recasting the familiar public goods strategies of punishers, cooperators, and defectors in this light. We then introduce a strategy—the informant—with no clear analog in public goods games and show that individuals employing this strategy are a key to the emergence of systems where cooperation dominates. We also find that a defection-dominated regime may be transitioned to one that is cooperation-dominated by converting an optimal number of players into informants. We discuss these findings, the role of informants, and possible intervention strategies in extreme adversarial societies, such as those marred by wars and insurgencies.

DOI: [10.1103/PhysRevE.82.066114](https://doi.org/10.1103/PhysRevE.82.066114)

PACS number(s): 89.65.Ef, 02.50.Le, 87.23.Ge

I. INTRODUCTION

The classic game of prisoners' dilemma is a simple model for social interactions in which selfishness and mistrust lead to a lack of cooperation among participants, even though such cooperation from all players would lead to the greatest collective payoff. However, despite the fact that many social interactions seem to conform to the payoff structure of the prisoners' dilemma, cooperation abounds in human activity. In public goods interactions, punishment, which decreases the payoff of defecting individuals, seems to provide a mechanism to explain this fact [1–3]. Yet, in many situations, punishment itself may be costly to the punisher, leading to the second-order “altruistic punishment” problem: how can cooperation proliferate if willingness to punish seems to be a necessary, but also costly, prerequisite? Again, in terms of public goods games, several solutions to this problem have been proposed; these include making the game optional [4], including both individual and group selections [5], and allowing for reputation-seeking behavior [6], among others [7–9].

Despite these advances, little work has focused on cooperation and punishment in more adversarial games where defecting players actively seek to harm others for their own benefit, rather than adopting the more passive free-rider role of the public goods game. Such games may be quite central to the understanding of many human activities not readily modeled by the public goods approach, such as warfare or crime. Furthermore, although adversarial games may present

many of the familiar conundrums discussed above, they may also admit new solutions.

Take, for example, criminal activity (which we shall use as the central example throughout this paper, although our results should not be limited to this domain). It is known that civilian actors play a critical role in creating and maintaining peaceful communities by forming communal relationships that lead to the acceptance of social norms within the group and to the shared understanding of what constitutes violation of such norms [10]. This leads in turn to self-regulating communities with informal safety and control measures [5,8], such as citizen-based surveillance and a positive disposition toward reporting offenses, helping to manage and discourage crime [11,12].

On the other hand, the temptation to violate social norms appears to be pervasive [7,6,13], and empirical evidence suggests that crime proliferates in highly disorganized societies [14] where the creation of a shared sense of social responsibility is hindered [15–18]. Furthermore, under extreme social circumstances where crime is highly prolific, the tendency to be concerned with the lawfulness of one's immediate neighborhood may be superseded by fear of retaliation for cooperating with authorities. This may happen even in social settings characterized by homogeneous populations or long-time residents. Fear may become so ingrained within a group that the accepted social norm is not to collaborate for the greater good, so that witnesses and even victims of crime choose not to cooperate with law enforcement in the prosecution of criminals. These behaviors generally emerge when forms of highly organized crime, such as the Italian Mafia [19], street gangs, or large drug cartels [20], come to permeate the social fabric [21].

*dorsogna@csun.edu

This contrast between low-crime (cooperation-dominated) and high-crime (defection-dominated) societies again highlights the importance of punishment in enhancing cooperation, but displays the same altruistic punishment problem discussed previously. Here, we attempt to solve this problem within an evolutionary adversarial game, recasting the familiar public goods strategies of punishing, cooperating, and defecting in relation to criminal behavior, while introducing a strategy, the informant, which will be shown to be critical to the eventual dominance of cooperation and punishment.

II. MODEL

We consider an idealized society where possible strategies vis-à-vis criminal behavior are divided into four categories across two distinct domains: those who either will or will not commit crimes and those who either will or will not report and serve as witnesses to crimes. Each individual thus possesses both a criminal trait and a reporting trait, so that there are four types of citizens (strategies): “paladins,” “villains,” “apathetics,” and “informants.” The P paladins represent model citizens who will never commit crimes and always cooperate with authorities, akin to public goods punishers. The V villains are active criminal offenders who do not cooperate with law enforcement, corresponding to defectors. In between these extremes are the A apathetic citizens who do not actively commit crimes, but who also choose to neither report offenses nor testify in cases where they are witnesses, perhaps due to inherent apathy or to fears of retaliation or ostracism. These apathetics are similar to second-order free riders [3,22] in other models; i.e., they cooperate at first order by not committing crimes, but defect at second order by not punishing offenders. Finally, the I informants are active criminals who will nevertheless cooperate with the justice system in the investigation of crimes they themselves did not commit, presumably because of personal gain [23] stemming either from a “*mors tua, vita mea*” attitude or from special enticements from the law enforcement system. They also represent a strategy with no clear analog in public goods games, in that they defect at first order by committing crimes, but cooperate at second order by punishing other criminals. In Fig. 1 we summarize the different types of citizens in our idealized society, which is assumed closed, so that the total population $N = P + V + A + I$ is a constant.

Our model then consists of a dynamic social game centered on criminal events and their aftermaths. The game unfolds iteratively in rounds by first determining the time t_c at which the next crime occurs (a Poisson arrival process) along with a random criminal player drawn from the $V + I$ pool of villains and informants who will be the victimizer for the said crime. Next, a random victim is chosen from the $N - 1$ remainder of the population; note that a victim may be of any citizen type, consistent with victimization surveys [24]. These two players each begin the game with a unitary payoff (without loss of generality), and the pair represents a single criminal act wherein the victimizer attempts to “steal” an amount $\delta \leq 1$ from the victim’s payoff. Each crime is thus associated with an immediate potential payoff loss δ for the victim and payoff gain δ for the victimizer.




FIG. 1. (Color online) Our model society is composed of four citizen types, defined by their propensities to both commit crimes and serve as witnesses in criminal investigations. Arrows between types indicate the number of possible game pairings and outcomes in which the strategy update step leads to a net change from one type to the other (e.g., there are two combinations that convert a villain into a paladin). Circular arrows within each type indicate those strategy updates that keep a player’s strategy unchanged within the given type.

The outcome of the criminal encounter depends on victim type. If the victim is either an apathetic or a villain, the crime is not reported to the authorities and is therefore successful: the victim’s payoff is decreased to $1 - \delta$ and the victimizer’s is increased to $1 + \delta$. If the victim is a paladin or an informant, the crime is reported to the authorities and an “investigation” begins. Here, an M member subsample of the $N - 2$ remaining individuals in the population is randomly selected as potential witnesses. This group consists of $m_P \leq P$ paladins and $m_I \leq I$ informants that serve as witnesses, along with $m_A \leq A$ apathetics and $m_V \leq V$ villains who do not. Hence, the witnessing fraction of the subsample is $w = (m_P + m_I) / M \leq 1$, a number used to determine the outcome of the investigation. With probability w , the victimizer is convicted: the victim’s payoff returns to its original unitary value, while the victimizer’s is decreased to $1 - \theta$, where $\theta \leq 1$ measures the severity of punishment. With probability $1 - w$, though, the crime is left unpunished: the victimizer’s payoff is increased to $1 + \delta$, while the victim’s is decreased to $1 - \delta - \epsilon$. The additional loss $\epsilon \leq 1 - \delta$ borne by the victim in this case may be interpreted as damages to his or her personal image or credibility, or as a loss of “faith in the system” after making an accusation that is unsubstantiated by the community. This additional loss may also be interpreted as revenge perpetrated by the accused who feels empowered by the lack of witnesses to the original crime; this may be especially important in extreme societies where retaliation is prevalent. The choice of reporting one’s victimization to authorities may thus be more detrimental than the original criminal act, especially in societies where few people will serve as witnesses to crimes.

The model is completed by specifying a method for citizens to change strategies over time; we have chosen a variant of the “proportional imitation” rule [25]. At the end of each




FIG. 2. (Color online) There are two qualitatively different equilibrium states available in our model, displayed here for both the [(a) and (b)] stochastic and [(c) and (d)] deterministic formulations. Paladins are shown as (red) circles, apathetics as (blue) squares, informants as (orange) triangles, and villains as (green) diamonds. All simulations use parameters $N=1000$, $M=5$, $\delta=0.3$, $\theta=0.6$, and $\epsilon=0.2$. [(a) and (c)] An initial population distribution $P=A=0$, $I=40$, $V=960$ evolves toward utopia, where only noncriminals remain. [(b) and (d)] An initial population distribution $P=400$, $A=I=0$, $V=600$ ends in dystopia, with only villains and apathetics remaining.

round, the player with the smaller payoff [26], denoted as the “loser,” performs a strategy update. Thus, we have assumed that our citizens possess a sort of “strategy inertia,” whereby they tend to retain their current strategy unless subject to a potentially traumatizing experience, such as being the victim of a crime that goes unpunished or being convicted of a crime. The loser selects his new strategy by choosing to emulate either the victim or the victimizer with probability proportional to each player’s payoff for that round. If the victimizer is emulated, the loser simply adopts the victimizer’s strategy and ends the update as either a villain or an informant. If the victim is emulated, the loser mimics the victim’s propensity to serve as a witness but, having sided with the victim of the crime, adopts a noncriminal strategy regardless of whether the victim was of a criminal type or not, ending the update as either a paladin or an apathetic. This removes an inherent bias of our game whereby each round is guaranteed to contain at least one criminal type (the victimizer), but is not guaranteed to contain a noncriminal type (since the victim may be an informant or villain). It also allows losing criminal victims to convert into noncriminal types; this may be thought of as a sort of “empathy” on the part of the loser, who may decide to not commit crimes in the future, not necessarily because this will maximize utility, but because he or she has experienced firsthand the consequences of being a victim.

III. RESULTS

For any given initial distribution of citizen types, our model may be iterated until one of two possible final equilibrium states is reached. The first equilibrium state is one that contains no criminals ($V+I=0$, $P+A=N$), denoted as the “utopian society.” Here, since no criminals are present, no further crimes nor further changes in strategies arise. An ex-

ample of a dynamical evolution of the system toward utopia is illustrated in Fig. 2(a). The other possible equilibrium outcome is that there are no citizens remaining that will cooperate with authorities ($P+I=0$ and $V+A=N$), denoted as the “dystopian society.” Here, while strategy updates continue to occur between villains and apathetics, these maintain the status quo with $V+A=N$ since there are no cooperative citizens to mimic. An example of a dynamical evolution leading to dystopia is illustrated in Fig. 2(b).

It is interesting to note that neither of these equilibrium states includes any informants. However, whether the system evolves toward utopia or dystopia is strongly dependent on the initial number of informants in the society. Figure 3 il-




FIG. 3. (Color online) Even a small number of informants substantially increase the likelihood that a society will evolve toward utopia. Here, we plot the probability of the stochastic system ending in utopia given the initial numbers of paladins (P_0) and informants (I_0), assuming zero initial apathetics; higher I_0 values lead to greater probability of ending in utopia. Simulations used $N=100$, $M=5$, $\delta=0.3$, $\theta=0.6$, and $\epsilon=0.2$, and were run 1000 times for each initial condition.

illustrates this effect for $N=100$ and initial conditions $P=P_0$, $I=I_0$, $A=0$, and $V=N-P_0-I_0$, displaying the fraction of simulations ending in utopia for varying initial conditions. In the case $I_0=0$, the society will never contain any informants, and utopia is never reached for values of P_0 less than some threshold; for the parameters of Fig. 3, this threshold is $P_0 \approx 35$. So, unless society is heavily populated with law abiding, cooperative citizens to begin with, a criminal-free steady state will not be found in the absence of informants.

However, if $P_0=34$, just below the threshold, and just one of the initial 66 villains is converted to an informant, so that $I_0 \rightarrow 1$ and $V_0 \rightarrow 65$, the system evolves toward utopia approximately 36% of the time. Hence, changing the second-order cooperativity of only a small fraction of the initial

criminals has a large positive impact on the final state of our society. More provocatively, if we leave the initial number of villains unchanged at $V_0=66$ and instead convert one of the 34 initial paladins into an informant (i.e., turn a second-order cooperating noncriminal into a second-order cooperating criminal), so that $I_0 \rightarrow 1$ and $P_0 \rightarrow 33$, the system again evolves toward utopia approximately 36% of the time; the old axiom “if you can’t beat ’em, join ’em” seems especially valid here. These qualitative trends persist under all parameter regimes.

An understanding of this behavior can be found by deriving a deterministic version of our stochastic model, which takes the form of four coupled ordinary differential equations (ODEs):

$$\dot{P} = (I+V) \left[(P+I)^2 \frac{1}{2-\theta} + I(A+V) \frac{1-\delta-\epsilon}{2-\epsilon} - P(A+V) \frac{1+\delta}{2-\epsilon} \right], \quad (1)$$

$$\dot{A} = (I+V) \left[V \frac{1-\delta}{2} - A \frac{1+\delta}{2} \right], \quad (2)$$

$$\dot{I} = I \left[(A+V) \frac{1+\delta}{2} + P(A+V) \frac{1+\delta}{2-\epsilon} - (P+I)^2 \frac{1}{2-\theta} - I(A+V) \frac{1-\delta-\epsilon}{2-\epsilon} - V(A+V) \right], \quad (3)$$

$$\dot{V} = V \left[(P+I)(A+V) \frac{1+\delta}{2-\epsilon} + (A-I) \frac{1+\delta}{2} - (P+I)^2 \frac{1}{2-\theta} - (I+V) \frac{1-\delta}{2} \right], \quad (4)$$

where P , A , I , and V are now understood to be fractions of the total population, so that $P+A+I+V=1$. Because of this conservative property, one can actually eliminate any one of the four equations above, and in practice we typically eliminate Eq. (2), leaving us with only P , V , and I . Output from these deterministic equations can be seen in Figs. 2(c) and 2(d), with initial conditions and parameters the same as those used in the stochastic model for Figs. 2(a) and 2(b); the stochastic and deterministic models are in agreement due to the relatively large value of N chosen (1000).

The value of I_0 strongly influences the dynamics of Eqs. (1)–(4). If $I_0=0$, I will be zero for all time (since $\dot{I}=0$), and Eq. (1) can be rewritten as

$$\dot{P} = \frac{1-A-P}{(2-\epsilon)(2-\theta)} \{ P^2 [4-\epsilon-\theta+\delta(2-\theta)] - P(2-\theta)(1+\delta) \}. \quad (5)$$

Equation (5) defines a threshold value P_s , given by

$$P_s \equiv \frac{(2-\theta)(1+\delta)}{4-\epsilon-\theta+\delta(2-\theta)}, \quad (6)$$

such that, if $P_0 < P_s$, then $\dot{P} < 0$ and P will asymptotically decay to zero. In this regime, the system evolves toward the dystopian fixed point, located at

$$A = A_d \equiv \frac{1-\delta}{2}, \quad V = V_d \equiv \frac{1+\delta}{2}. \quad (7)$$

For the parameters of Figs. 2(b) and 2(d), this fixed point is at $A_d=0.35$ and $V_d=0.65$, matching the observed final proportions of villains and apathetics in the stochastic simulations. If, on the other hand, $P_0 > P_s$, then $\dot{P} > 0$ and P will increase until $A+P=1$, the utopian society fixed line. Hence, when $I_0=I=0$, the trajectory with $P=P_s$ serves to partition the P - V phase plane into the two basins of attraction for utopia and dystopia. On the $P=P_s$ trajectory there is a final fixed point characterized by

$$A = A_s \equiv \frac{(2-\epsilon)(1-\delta)}{2[4-\epsilon-\theta+\delta(2-\theta)]}, \quad (8)$$

$$V = V_s \equiv \frac{(2-\epsilon)(1+\delta)}{2[4-\epsilon-\theta+\delta(2-\theta)]}.$$

Analysis shows that this is a saddle point, explaining why it is not observed in the stochastic simulations. For the parameters used in Fig. 3, $P_s \approx 0.5$, and at $I_0=0$ and $P_0=0.5N=50$ the system evolves to utopia for approximately half of the stochastic simulations.




FIG. 4. (Color online) All trajectories with $I_0 > 0$ end in utopia in the deterministic system. This example phase portrait shows unstable fixed points in light gray (red), unstable fixed lines in thick light gray (red), stable fixed lines in thick dark gray (blue), and trajectories beginning (or ending) along various eigenvectors as thick medium gray (green) arrows. Note that the dystopian fixed point d and saddle point s are unstable to increases in I , so that the only attracting final states for $I_0 > 0$ are those utopias with $P > P_c$. Here, $\delta=0.3$, $\theta=0.6$, and $\epsilon=0.2$, but the qualitative results are independent of parameters.

The situation is quite different when $I_0 > 0$ (refer to Fig. 4). Each of the three fixed points (or lines) described above possesses two eigenvectors that lie in the P - V plane (we refer to these as $\mathbf{v}_1$ and $\mathbf{v}_2$) and one that has a nonzero I component (we refer to this as $\mathbf{v}_3$). For the saddle point at $(P_s, A_s, 0, V_s)$, $\mathbf{v}_1$, which corresponds to the trajectory $P=P_s$, has a negative eigenvalue, while $\mathbf{v}_2$ and $\mathbf{v}_3$ both have positive eigenvalues. For the utopian fixed line $P+A=1$, $\mathbf{v}_1$, which corresponds to the P axis, has a zero eigenvalue since any given utopia is neutrally stable with respect to all other utopias. The other two utopian eigenvectors, $\mathbf{v}_2$ and $\mathbf{v}_3$, share an eigenvalue that is negative for $P > P_c$ where

$$P_c \equiv P_s \left(1 - \frac{\epsilon}{4} \right) \sqrt{1 + \frac{8(2-\epsilon)^2}{(1+\delta)(2-\theta)(4-\epsilon)^2}} + \frac{P_s \epsilon}{4}, \quad (9)$$

and positive for $P < P_c$; hence, utopian states with $P > P_c$ are attractors to at least some volume of phase space, including some states with $I > 0$. The dystopian fixed point has negative eigenvalues for both $\mathbf{v}_1$, which corresponds to the V axis, and $\mathbf{v}_2$, but has a zero eigenvalue for $\mathbf{v}_3$, making a full classification of the point somewhat difficult. However, upon diagonalizing the system via the dystopian eigenvectors, one finds that the first nonzero term in the dynamics along $\mathbf{v}_3$ is positive, meaning that the dystopian fixed point is weakly unstable to the addition of informants. Therefore, the only attractor for initial conditions with nonzero I is utopia (specifically with $P > P_c$), and the deterministic system will *always* end in utopia if $I_0 > 0$.

On the other hand, the time to reach utopia may be quite long, and if I_0 is extremely small (and $P_0 \leq P_s$) the system may first spend a large amount of time very near the dystopian fixed point due to the zero eigenvalue there. Our deterministic simulations match this behavior, always ending in utopia when $I_0 > 0$ regardless of the parameters or initial conditions we choose, although sometimes spending very

large amounts of time near dystopia first. This is in contrast to the stochastic system, which does not always end in utopia for $I_0 > 0$, as evidenced in Fig. 3. This difference arises mainly because random fluctuations in the stochastic system may cause an initially small population of informants to die out during the time spent near dystopia. We expect the stochastic annihilation of informants to be more likely for smaller values of I_0 , which is also seen in Fig. 3.

Because of the large influence that informants have on determining the final state of the system, one possible strategy for undermining a dystopian society is to convert some number of villains (or apathetics) into informants, which should lead to utopia with a probability dependent on how many conversions took place (see Fig. 3). Presumably, however, prospective converts will require an incentive such as monetary payment or protection guarantees to switch their strategy to informant, since cooperating with authorities in a dystopian society is inherently risky. We denote this incentive cost to convert a single villain to an informant as α . A low number of converts keep the total incentive costs low, but too few converts may result in an unreasonably long period of time before utopia is reached, if it is reached at all, leading to very high total losses borne by victims of crime, denoted as L_{vic} , in the process. Additionally, too few converts may leave the system in dystopia, necessitating a number of conversion rounds R greater than 1. Thus, we expect that an optimal number of conversions I_0^* may exist that minimize the total cost of reforming the society,

$$C = R\alpha I_0 + L_{\text{vic}}. \quad (10)$$

Plots of C versus I_0 for both the deterministic system and the average of 1000 stochastic simulations are shown in Fig. 5. Both the stochastic and deterministic formulations indeed exhibit an optimal I_0^* in this case, although the existence and value of I_0^* depend on the specific parameters used. Our simulations suggest that I_0^* exists for a large range of parameter values and, when it does not exist, costs are simply a decreasing function of I_0 . The stochastic and deterministic costs agree quite well at high I_0 , where the stochastic system is very likely to evolve to utopia after just one round of conversions (see Fig. 3), as the deterministic system does. At very low I_0 , the stochastic system displays a much lower cost than the deterministic one, despite the fact that R for the stochastic system is likely greater than 1 in this regime. This is because the stochastic system will tend to spend less time very near the dystopian state than the deterministic system in this case. The stochastic system will instead reach dystopia and trigger another round of conversions, whereas the dynamics of the deterministic system cause it to linger in close proximity to the dystopian fixed point, greatly increasing the costs borne by victims.

IV. DISCUSSION

In this work we have analyzed the evolution of cooperative behavior within extreme adversarial societies. We introduced a game strategy, the informant, for which no analog exists within evolutionary public goods games, showing that this figure may dramatically alter game dynamics.




FIG. 5. (Color online) An optimal number of villain-to-informant conversions I_0^* may be made to undermine a dystopian society while minimizing reformation costs. The solid line and points are the averaged simulated costs C versus I_0 for 1000 runs of the stochastic system, while the dashed (red) line is the result for the deterministic system. Both systems use parameters $N=100$, $M=5$, $\delta=0.3$, $\theta=0.6$, $\epsilon=0.2$, and $\alpha=20$.

This research has direct implications for the design of intervention strategies in extreme security settings where few members of the general populace are willing to cooperate with the law enforcement, military, or judicial systems to reign in crime and/or violence. Fear of reprisal, perception that the security and legal systems are ineffective, or even a culture of noncooperation with authority may prevent individuals from coming forward to report and participate in the prosecution of crime. We note that such conditions are particularly prevalent in areas experiencing an active insurgency such as Afghanistan. Our results suggest that if a community is marred by extreme levels of crime or violence and contains only insurgents (criminals), model citizens, and passive citizens, then it is very unlikely if not impossible to dislodge the community from this state simply by trying to eliminate the insurgents. Our observation is consistent with recent reevaluations of the dynamics of insurgencies [27,28]. We have

shown, though, that introducing even a single individual who is willing to inform when they or others are victimized, while not desisting from committing crimes themselves, can undermine criminal and violent activity and ultimately lead to its eradication. Cultivating larger numbers of informants speeds up this process and increases the likelihood of its success. Our results also suggest that there may be an optimal number of informants that reduces the overall costs of converting a dire security situation into a manageable one.

Future work will explore how the presence of mixed strategies affects our findings. The current model allows individuals to choose from only four very extreme pure strategies, each of which behaves in completely predictable ways (e.g., paladins never commit crimes and always cooperate with authorities). However, a more realistic model might allow each individual to act in a stochastic manner characterized by two probabilities: the probability p_c of committing criminal acts and the probability p_w of serving as a witness. Each of these probabilities could then take on any value between 0 and 1, changing in prescribed ways as interactions took place. In the continuum sense, one would then focus on the dynamics of the macroscopic probability distributions of the variables p_c and p_w . The current model treats these distributions as delta functions located at $p_c=V+I$ and $p_w=P+I$, but a more general model would allow these distributions to take on other forms. This would necessitate the switch from an ODE formulation of the model [Eqs. (1)–(4)] to an integro-differential form [29], possibly leading to new behavior.

ACKNOWLEDGMENTS

This work was supported in part by the U.S. National Science Foundation Mathematics (Grants No. DMS-0719462, DMS-0968309, DMS-0907931, and DMS-1021850), Office of Naval Research (Grant No. N000141010221), Air Force Office of Scientific Research (Grant No. FA9550-10-1-0569), and Army Research Office (Grants No. W911NS-09-1-0559 and No. 50363-MA-MUR).

- [1] J. Henrich and R. Boyd, *J. Theor. Biol.* **208**, 79 (2001).
- [2] R. Boyd, H. Gintis, and S. Bowles, *Science* **328**, 617 (2010).
- [3] K. Panchanathan and R. Boyd, *Nature (London)* **432**, 499 (2004).
- [4] C. Hauert, A. Traulsen, H. Brandt, M. Nowak, and K. Sigmund, *Science* **316**, 1905 (2007).
- [5] R. Boyd, H. Gintis, S. Bowles, and P. Richerson, *Proc. Natl. Acad. Sci. U.S.A.* **100**, 3531 (2003).
- [6] E. Fehr and U. Fischbacher, *Nature (London)* **425**, 785 (2003).
- [7] E. Ostrom, *J. Econ. Perspect.* **14**, 137 (2000).
- [8] E. Fehr and U. Fischbacher, *Evol. Hum. Behav.* **25**, 63 (2004).
- [9] M. Janssen, R. Holahan, A. Lee, and E. Ostrom, *Science* **328**, 613 (2010).
- [10] W. Skogan, *Crim. Justice* **10**, 39 (1988).
- [11] R. Sampson, S. Raudenbush, and F. Earls, *Science* **277**, 918 (1997).
- [12] R. Bursik and H. Grasmick, *Neighborhoods and Crime: Dimensions of Effective Community Control* (Lexington Books, New York, 1993).
- [13] L. Cohen and R. Machalek, *Am. J. Sociol.* **94**, 465 (1988).
- [14] W. Skogan, *Disorder and Decline: Crime and the Spiral of Decay in American Neighborhoods* (University of California Press, Los Angeles, 1990).
- [15] G. Jensen, *The Encyclopedia of Criminology: Social Disorganization Theory* (Routledge, New York, 2005).
- [16] R. Sampson and W. Groves, *Am. J. Sociol.* **94**, 774 (1989).
- [17] J. Conklin, *Criminology* (Pearson Education, Boston, 2007).
- [18] R. J. Sampson, *Crim. Justice* **8**, 271 (1986).
- [19] D. Gambetta, *Trust: Making and Breaking Cooperative Relations* (Blackwell, Oxford, 1988).
- [20] J. Beittel, *Mexico's Drug Related Violence* (Congressional Research Service, Washington, D.C., 2009).
- [21] S. Mallory, *Understanding Organized Crime* (Jones and Bartlett Publishers International, London, 2007).

- [22] E. Fehr and S. Gächter, *Nature (London)* **415**, 137 (2002).
- [23] R. Rosenfeld, B. Jacobs, and R. Wright, *Brit. J. Criminol.* **43**, 291 (2003).
- [24] R. Rosenfeld, J. Lynch, and L. Addington, *Understanding Crime Statistics: Revisiting the Divergence of the NCVS and the UCR* (Cambridge University Press, Cambridge, England, 2006).
- [25] K. Schlag, *J. Econ. Theory* **78**, 130 (1998).
- [26] For some limiting cases, such as $\theta=0$, the payoff of both the victimizer and victim may be equal. In this case, one of them may be chosen at random to represent the loser, although we do not explore these cases here.
- [27] J. Nagl and P. Schoomaker, *Learning to Eat Soup with a Knife: Counterinsurgency Lessons from Malaya and Vietnam* (University of Chicago Press, Chicago, 2005).
- [28] D. Kilcullen, *The Accidental Guerrilla: Fighting Small Wars in the Midst of a Big One* (Oxford University Press, Oxford, 2009).
- [29] A. Boccabella, R. Natalini, and L. Pareschi (unpublished).