Euro. Jnl of Applied Mathematics (2013), *vol.* 24, *pp.* 131–159. © Cambridge University Press 2012
doi:10.1017/S0956792512000332

External conversions of player strategy in an evolutionary game: A cost-benefit analysis through optimal control

MARTIN B. SHORT[1], ASHLEY B. PITCHER[2] and MARIA R. D'ORSOGNA[3]

[1]*Mathematics Department, UCLA, Los Angeles, CA 90095-1555, USA*
email: mbshort@math.ucla.edu
[2]*Centre d'Analyse et de Mathématique Sociales, Ecole des Hautes Etudes en Sciences Sociales, 190-198 Avenue de France, Paris cedex 13, France 75244*
email: ashley.pitcher@balliol.oxon.org
[3]*Mathematics Department, CSUN, Los Angeles, CA 91330-8313, USA*
email: dorsogna@csun.edu

(*Received 27 October 2011; revised 17 September 2012; accepted 17 September 2012; first published online 10 October 2012*)

We consider an optimal control problem based on the evolutionary game theory model introduced by Short *et al.* (Short, M. B., Brantingham, P. J. & D'Orsogna, M. R. (2010) Cooperation and punishment in an adversarial game: How defectors pave the way to a peaceful society. *Phys. Rev. E* **82**(6), 066114.1–066114.7) to study societal attitudes in relation to committing and reporting crimes. Since in [26] (Short, M. B., Brantingham, P. J. & D'Orsogna, M. R. Cooperation and punishment in an adversarial game: How defectors pave the way to a peaceful society. *Phys. Rev. E* **82**(6), 066114.1–066114.7) it is shown that the presence of criminal informants leads to diminishing crime, in this paper we investigate the active recruitment of informants from the general population via external intervention, albeit at a cost to society. While higher recruitment levels may be the most beneficial in abating crime, these are also more expensive. We thus formulate our optimal control problem to account for finite resources, incurred costs and expected benefits, and determine the most favourable recruitment strategy under given constraints. We consider the cases of targeted and untargeted recruitment, and allow recruitment costs to depend on past cumulative payoffs within a given memory time-frame so that conversion of more successful individuals may be more costly than that of less successful ones. Our optimal control problem is expressed via three control functions subject to a system of delay differential equations, and is numerically solved, analysed and discussed under different settings and in different parameter regimes. We find that the optimal strategy can change drastically and abruptly as parameters and resource constraints vary, and that increased information on individual player strategies leads to only slightly decreased minimal costs.

Key words: Optimal control, Evolutionary Game theory, Sociological modeling, Crime

1 Introduction

The study of social phenomena using mathematical and statistical physics methods has attracted much interest in recent years [11]. Opinion formation [3, 20, 28], voter

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

models [19, 27], racial segregation [13, 30, 31], linguistic changes [1, 5, 16] and Internet and social network dynamics [7, 21, 22] are just a few of the many areas that have been analysed within quantitative and mathematical frameworks. In some cases, the findings derived from these models have been useful in predicting outcomes of social policy and towards planning real-life strategies.

One of the topics that has received considerable attention within the broad field of social sciences is human cooperation [10, 25], especially in the context of the so-called 'prisoner's dilemma' scenarios [24] where each individual enjoys selfish gains by defecting from mutual collaborations, but society as a whole benefits if all members work together for the greater good. Typically, the goal of prisoner's dilemma studies is to understand how social cooperation may arise if incentives, inequalities or asymmetries between players are introduced [14, 17].

Intimately connected to the study of social cooperation is the study of criminal behaviour. Criminal acts may be interpreted as a form of defection, where the criminal chooses to pursue his or her personal gain over the common societal good of abiding by the law. If the expected benefits derived from committing crimes – such as monetary gain or increased social standing – exceed the expected losses – such as punishment or social stigma – then criminal behaviour may present a form of prisoner's dilemma. Within this context, the role of law enforcement and government agencies is to determine practices that will reduce and prevent 'criminal-defection' acts, allowing the 'crime-free collective' interest to prevail.

The question of how to best intervene to build a law-abiding society has long been the subject of discussion among criminologists, sociologists, urban planners and policy makers [4], especially for high-crime locations where abundant opportunities and rewards for illegal behaviour far outweigh the pursuit of communal interests. Aside from traditional methods used as deterrents or in response to crime, such as incarceration and increasingly harsh punishments, other avenues have begun to attract interest for being more beneficial and cost-effective. These include better social, educational and employment policies to prevent crimes, and better communication between law enforcement and residents through programmes designed to create trust and foster engagement [29].

As with most social issues, there are several ways of allocating a finite pool of resources, and finding the optimal way to intervene against the proliferation of crime may depend on the particular social fabric at hand. Programmes based on coordinated efforts between residents and local agencies may not always yield desired results due to abuse or mistrust. For example, in communities dominated by large-scale organized crime or powerful warlords, such as the Mafia, street gangs, large drug cartels or the Taliban, the acceptance of crime may be ingrained in the culture. Sometimes criminals may also be revered by the community, especially if the criminal network is deeply rooted within the society, providing residents with unorthodox services, protection and employment. In these cases, fears of retaliation, mistrust of authorities and loyalty to gangs or perpetrators may render any direct attempt to foster cooperation with law enforcement officials extremely difficult.

The same cultural mistrust that prevents active partnerships between government officials and citizens may also prevent law enforcement from gathering evidence against criminals due to lack of leads or witnesses willing to testify when crimes occur. In these

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

cases, the presence of an informant, who provides information while remaining within the active criminal milieu, may be of pivotal importance in solving crimes and, from a larger viewpoint, in helping to foster a climate of change. If informants lead to the arrest of criminals, a deterrence feedback loop may be created where residents favour a more cooperative approach with authorities, away from criminal associations.

The benefits of having access to informant figures may, however, come with a price. Criminal witnesses to crimes may only be persuaded to yield information in exchange for protection, money or lighter sentences after an arrest. Likewise, authorities may employ undercover agents who remain active for long periods of time and who may even be called upon to commit criminal acts. In all these cases, luring a criminal, or training any type of citizen to become an informant, requires the expenditure of resources via direct monetary compensation or via the institution of specific police or community programmes. Available resources must thus be divided among punishment, prevention and informant recruitment efforts in the most efficient way.

In previous work [26], we introduced an evolutionary game-theoretic model for a society where each citizen may or may not commit crimes and may or may not cooperate with authorities. Within this context, one possible citizen strategy is to be a criminal informant, one who commits crimes but also shares information with law enforcement. The dynamics of our model society depends on how criminal acts unfold between players, on how citizens alter their strategies and on several parameters, including the level of punishment for criminal offenses, the severity of retaliation against witnesses to crimes and the expected benefits for successful criminals. We found that the initial number of informants has a dramatic effect on the final societal configuration. Indeed, within our idealized model, informants are able to eventually drive society towards stable and crime-free configurations that would not have been achieved in their absence.

The goal of this paper is to utilize optimal control theory to determine the most favourable recruitment rates of informants from the general population, within the model presented above and assuming certain forms for the recruitment and carrying costs. Recruitment costs depend on the number of citizens converted to the informant role, while carrying costs instead represent the price society must pay for the occurrence of crimes before a low crime state is reached. We consider two possibilities for the recruitment process: targeted recruitment, where authorities have prior knowledge of citizen strategies; and untargeted recruitment, where authorities lack this information. Untargeted recruitment is associated with the presence of a memory kernel as discussed later in the text. We compare and contrast results from the two possible conversion scenarios in terms of the cost to recruit single individuals, the parameters of the model and the maximal recruitment rate – a proxy for the total resource pool the authorities have at their disposal to convert citizens to informants. We find that the optimal recruitment policy may change drastically and discontinuously as these parameters are varied, and that targeted recruitment offers only slight benefits over untargeted recruitment in terms of reduced minimal costs. These results highlight the difficulty in determining the optimal course of action in any real-world scenario, where information about parameters and current system state may be highly limited, and also call into question the benefits of expending resources to obtain highly detailed knowledge of individual citizens' current strategies in order to implement a more targeted form of recruitment.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

Our paper is organized as follows. In Section 2 we introduce our previous model, describing our idealized society and its time evolution. In Section 3 we introduce and define the optimal control problem. In Section 4 we find and discuss solutions to the optimal control problem in a variety of cases. Finally, in Section 5 we discuss our results and make some final concluding remarks.

2 The model

To study the dynamics of citizen attitudes towards committing crimes and cooperating with authorities, we use the game-theoretic model described in detail in [26]. In brief, the model assumes an idealized society composed of four types of citizens: (1) paladins, who do not commit crimes and always cooperate with authorities; (2) apathetics, who do not commit crimes but also do not cooperate with authorities; (3) villains, who commit crimes and do not cooperate with authorities and (4) informants, who commit crimes but also cooperate with authorities.

The game unfolds through a succession of rounds, each characterized by the occurrence of a criminal act. At the beginning of each round, all players are assigned a utility payoff of unitary value. A victimizer is chosen at random among villains and informants (i.e. those players who will commit crimes), while the victim is chosen at random from the entire population, and may therefore be of any type. The criminal act between the two chosen players involves the victimizer 'stealing' an amount $\delta < 1$ from the victim's unitary payoff, thus adding it to his or her own payoff, which becomes $1 + \delta$.

If the victim of the crime is a type that does not cooperate with authorities (apathetic or villain), then the crime is automatically successful, as the authorities are not informed. On the other hand, if the victim is a type that will cooperate with authorities (paladin or informant), the crime is reported, and an 'investigation' begins. The latter involves the selection of a random group of citizens, each of whom either will (paladins or informants) or will not (apathetics or villains) serve as a witness to the crime. The probability that the victimizer is convicted for the crime is then equal to the fraction of witnesses that cooperate with authorities.

If the criminal is convicted, the victim's payoff is restored to unity, while the victimizer's is lowered by an amount $\theta \leqslant 1$ that represents the severity of punishment. However, if the criminal is not convicted, he or she keeps the δ taken from the victim, while the victim's payout is lowered by an additional amount $\epsilon \leqslant 1 - \delta$. The quantity ϵ represents retaliation against the victim for reporting crimes to the authorities. It is only a loss to the victim and is not transferred to the victimizer.

The model is completed by providing a method for citizens to change their strategies over time. At the end of each criminal act, the victimizer and victim each have their respective payoff. That player with the lower payoff is denoted as the 'loser', and is given the opportunity to change his strategy. To do so, he first chooses one of the two players at random with probability proportional to that player's payoff; we shall call this chosen player the 'target' (note that the loser and target can be the same person). The loser then adopts (or retains) the target's attitude towards cooperating with authorities. Furthermore, if the target was the victimizer for that game, then the loser becomes (or remains) a criminal type; if the target was the victim for that game, then the loser becomes (or remains) a non-criminal type.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

Our model may be cast either in a stochastic or a deterministic form, both of which have been thoroughly explored in [26]. The deterministic choice leads to the following set of coupled, nonlinear Ordinary Differential Equations (ODEs):

$$\dot{P} = f_P \equiv (1-P-A)\left[(1-A-V)^2\frac{1}{2-\theta} + (1-P-A-V)(A+V)\frac{1-\delta-\epsilon}{2-\epsilon} - P(A+V)\frac{1+\delta}{2-\epsilon}\right], \tag{2.1}$$

$$\dot{A} = f_A \equiv (1-P-A)\left[V\frac{1-\delta}{2} - A\frac{1+\delta}{2}\right], \tag{2.2}$$

$$\dot{V} = f_V \equiv V\left[(1-P-A)(A+V)\frac{1+\delta}{2-\epsilon} + (2A-P-V-1)\frac{1+\delta}{2} - (1-A-V)^2\frac{1}{2-\theta} - (1-P-A)\frac{1-\delta}{2}\right], \tag{2.3}$$

where P, A and V represent the fraction of paladins, apathetics and villains in the total population respectively. The fraction of informants I has been eliminated from the system using the constraint $P + A + V + I = 1$. In writing equations (2.1)–(2.3), we have rescaled time by the average rate ω at which individual criminals offend: the above equations are thus expressed in terms of a dimensionless time unit.

The derivation of these ODEs is illustrated in detail in our previous work [26]. As an example of how these are derived, though, consider equation (2.2). The number of apathetics in the society decreases when an apathetic is victimized, occurring at rate $(I + V)A$, and then chooses to mimic the criminal, which happens with probability $(1+\delta)/2$ (the payoff of the criminal normalized by the combined payoffs of both players). The number of apathetics may also increase after a villain is victimized, occurring at rate $(I + V)V$, and then chooses to 'mimic' the victim, which happens with probability $(1-\delta)/2$. The derivations of equations (2.1) and (2.3) follow similar arguments.

The functions f_P, f_A and f_V have been introduced to lighten the notation, since we will later refer to the right-hand sides of equations (2.1)–(2.3). The subscripts P, A and V indicate paladin, apathetic and villain strategies respectively; later, we will introduce other quantities that depend on strategy type, using the same subscript notation. Equations (2.1)–(2.3) are the starting point of this work, from which we will construct a total cost functional to be minimized.

As shown in [26], the long-term behaviour of the system of equations (2.1)–(2.3) is highly dependent on the initial number of informants, I_0. If $I_0 = 0$, then no informants will ever exist, and the system evolves to one of two qualitatively different states, depending on the initial number of paladins, P_0. If $P_0 < P_s$, where

$$P_s \equiv \frac{(2-\theta)(1+\delta)}{4-\epsilon-\theta+\delta(2-\theta)}, \tag{2.4}$$

then the system evolves to a point we call 'dystopia', located at

$$P = I = 0,\ A = A_d \equiv \frac{1-\delta}{2},\ V = V_d \equiv \frac{1+\delta}{2}\ . \tag{2.5}$$

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

In dystopia, there are no paladins or informants left, no one cooperates with authorities, most citizens are criminals, and all crimes are successful. If, on the other hand, $P_0 > P_s$, then the system evolves to one of a continuum of 'utopia' states, in which $V = I = 0$, $P + A = 1$ and $P > P_c$, where

$$P_c \equiv P_s \left(1 - \frac{\epsilon}{4}\right) \sqrt{1 + \frac{8(2-\epsilon)^2}{(1+\delta)(2-\theta)(4-\epsilon)^2}} + \frac{P_s \epsilon}{4}. \tag{2.6}$$

In utopia, there are no criminals left and no further victimizations occur. From a societal viewpoint, any of the utopia equilibrium points are preferred over dystopia.

For the case of a non-zero number of initial informants ($I_0 > 0$) one can show that the system of equations (2.1)–(2.3) will always end in a utopian state, since the dystopian fixed point is unstable to the addition of informants. However, this instability is algebraic, rather than exponential, so the system evolves rather slowly when near the dystopian state. Hence, though the inclusion of informants will always lead to utopia eventually, it may take a very long time for this to happen, with the society lingering in a near-dystopia, high-crime state in the meantime. For a system starting near dystopia with $I_0 \ll 1$, the approximate time to reach utopia, T_u, can be shown to be

$$T_u \approx \frac{1}{I_0(1-\epsilon/2)} \,. \tag{2.7}$$

Since a dystopian state may be 'reformed' towards utopia solely through the introduction of a small number of informants, a useful strategy towards a crime-free society may be that of actively recruiting informants. Doing so, however, comes at a price. In this paper, we consider two forms of costs when shifting from dystopia to utopia: recruitment costs, which are necessary to convert paladins, apathetics or villains into informants; and carrying costs, which are the crime-related losses suffered by citizens before society reaches utopia. One can immediately see that these two costs are in opposition to one another. On one hand, the authorities may choose to convert a very large number of citizens into informants, thus bearing a large recruitment cost. In this case, the system will evolve to utopia rather quickly so that the carrying costs are relatively small. On the other hand, if only a small number of citizens are converted to informants, recruitment costs may be kept low but it may take a very long time for society to reach utopia. In this case, carrying costs are quite large. Ideally, the recruitment rates of citizens will be chosen in an optimal way, to minimize the total costs of reforming society. We turn to this problem now.

3 The optimal control problem

Within our model, individuals may change into informants at the end of each round spontaneously due to the rules of the game. Here we wish to introduce the additional possibility of informant recruitment through external intervention, for example via specifically designed programmes that, by necessity, carry a cost to society. In particular, we consider two possible mechanisms of informant recruitment: targeted recruitment and untargeted recruitment.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

In targeted recruitment, we assume that authorities have prior knowledge of the strategies of all N individuals within the population so that recruitment is extended only to those whose profile is most desirable at a given point in time. We assume that the total recruitment rate is limited to a maximum value M, and that the dynamic control variables of our problem, $u_X(t)$, represent the fraction of M devoted to targeting citizens of strategy X, where $X \in \{P,\ A,\ V\}$. Hence, the controls are constrained so that

$$u_X \geqslant 0,\ u_P + u_A + u_V \leqslant 1. \tag{3.1}$$

Conversely, in untargeted recruitment we assume that authorities *randomly* choose potential recruits at a rate M from the general population, without knowing in advance the strategy of a citizen. Once a potential recruit is selected, his or her current strategy is revealed, and authorities may then choose whether or not to convert him or her into an informant. Given that the candidate's current strategy is X, where $X \in \{P,\ A,\ V\}$, the dynamic control variable $u_X(t)$ is defined as the probability that the authorities choose to enlist this individual. Since these are probabilities, these control variables are bounded at all times via

$$0 \leqslant u_X(t) \leqslant 1. \tag{3.2}$$

Note the slight difference between the interpretation of u_X in the targeted and untargeted cases due to the different assumptions made on the recruitment process.

Given these forms of external informant recruitment, equations (2.1)–(2.3) must be modified so that

$$\dot{X} = f_X - m \begin{cases} u_X & \text{for targeted recruitment} \\ X u_X & \text{for untargeted recruitment,} \end{cases} \tag{3.3}$$

where the f_X functions have been defined in equations (2.1)–(2.3) and $m = M/\omega N$ is the dimensionless form of our maximal recruitment rate. The quantity m can be interpreted as follows: Since ω is the average rate at which a criminal offends, N is the total size of the population, and M is the maximum recruitment rate, m represents the maximum fraction of the population that could be converted during the typical waiting time between a given criminal's offences. This is equivalent to the ratio of the typical time between crimes committed by an individual, $1/\omega$, and the typical time to convert all citizens, N/M. It may therefore be assumed that, realistically, $m \ll 1$, though we explore higher values of m in the following sections.

As stated previously, the active enlistment of informants carries a cost per recruited individual, which we describe by the functions $C_X(t, \tau)$, where $X \in \{P,\ A,\ V\}$ is again the current strategy of the individual being converted. We assume that the cost functions $C_X(t, \tau)$ depend on the expected average prior loss of individuals of strategy X over the previous τ time units so that the greater loss an individual has experienced within the $[t - \tau, t]$ interval, the less costly it will be to convert him or her to an informant at time t. In other words, if individuals of strategy X have 'faired poorly' and have experienced a large average loss over the period τ, they will be willing to become informants for a small price, since their current strategy has proven to be unsatisfactory. Conversely, if a particular citizen group is 'fairing well' and has had a small average loss or even an

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

average gain over τ, more incentives will be required to enlist individuals as informants, since their current adopted strategy is satisfactory. We thus choose cost functions of the form

$$C_X(t,\tau) = c\,[2 - \ell_X(t,\tau)]\,, \tag{3.4}$$

where c is a positive constant, and the $\ell_X(t,\tau)$ functions are the expected average prior loss functions over τ for citizens of strategy X. In constructing the $\ell_X(t,\tau)$ functions, we will use the convention that losses are included as positive terms and gains as negative ones. We note that one could define separate coefficients c_X for the various strategy types, but we do not do so here.

The choice of the memory length τ used in the prior loss functions will depend on the type of recruitment employed by authorities. In the case of targeted recruitment, we let authorities have general knowledge of citizen strategies, so we also assume that this knowledge is available to the general population. Hence, when determining their recruitment cost, these individuals are able to give a precise estimate of their current *instantaneous* rate of losses so that $\tau = 0$. On the other hand, in the case of untargeted recruitment, authorities have no information about the strategy of individuals until they select them. Hence, we assume that individuals in the general population also lack such broad knowledge. Therefore, when deciding the price they will require to be converted to informants, these individuals must rely on their past experiences so that τ should be strictly greater than zero.

Let us now construct the paladin loss function $\ell_P(t,\tau)$. Paladins never achieve any net gains, and experience a $\delta + \epsilon$ net loss only when they are victims of an unpunished crime. Here δ is due to the initial victimization and ϵ is the retaliation they receive for reporting the crime to the authorities. Since the victimization rate of an individual is $I + V$ and the probability that a crime against a paladin goes unpunished is $A + V$ (the fraction of witnesses that will not cooperate with authorities in the investigation), the loss function for paladins is

$$\ell_P(t,\tau) = \frac{1}{\tau}\int_{t-\tau}^{t} (I+V)(A+V)(\delta+\epsilon)\,ds \geqslant 0. \tag{3.5}$$

Note that since $I + V \leqslant 1$, $A + V \leqslant 1$ and $\delta + \epsilon \leqslant 1$, the average expected loss $\ell_P(t,\tau) \leqslant 1$. Hence, $C_P(t,\tau) \geqslant 0$.

We can similarly construct the loss function $\ell_A(t,\tau)$ for apathetic citizens by noting that, since they do not report to authorities, they always experience δ losses upon victimization. Thus, we find

$$\ell_A(t,\tau) = \frac{1}{\tau}\int_{t-\tau}^{t} (I+V)\delta\,ds \geqslant 0. \tag{3.6}$$

As with paladins, $\ell_A(t,\tau) \leqslant 1$ so that $C_A(t,\tau) \geqslant 0$.

Unlike paladins and apathetics, villains may experience both losses and gains so that $\ell_V(t,\tau)$ is not always positive. Losses to villains come in two forms: from victimization, which are the same as those for apathetics, and from punishment for failed crimes, which have magnitude θ. Punishment occurs when villains target paladins or informants, with probability $P + I$, and are convicted for their crimes, which also happens with probability $P + I$ (the fraction of witnesses who will cooperate in investigations). Let us call the

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

strictly construed villain losses $\ell_V^-(t,\tau)$, which are given by

$$\ell_V^-(t,\tau) = \frac{1}{\tau}\int_{t-\tau}^{t} (I+V)\delta + (P+I)^2\theta \, ds \geqslant 0. \tag{3.7}$$

Villains may also gain from their illicit activities in two different ways. If the villain victimizes an apathetic or another villain, which occurs with probability $A+V$, the crime is automatically successful and the villain gains δ. If the villain victimizes an informant or paladin, which occurs with probability $P+I$, the villain gains δ only if the crime is successful, which happens with probability $A+V$. Hence, the strictly construed villain gains $\ell_V^+(t,\tau)$ may be written as

$$\ell_V^+(t,\tau) = \frac{1}{\tau}\int_{t-\tau}^{t} (1+P+I)(A+V)\delta \, ds \geqslant 0. \tag{3.8}$$

Since gains are negative losses, the net loss function $\ell_V(t,\tau)$ for villains is given by

$$\ell_V(t,\tau) = \ell_V^-(t,\tau) - \ell_V^+(t,\tau) = \frac{1}{\tau}\int_{t-\tau}^{t} \delta\,[I - A - (P+I)(A+V)] + \theta(P+I)^2 \, ds. \tag{3.9}$$

Note that the maximum value for ℓ_V is given by $\delta+\theta \leqslant 2$, when $I=1$; hence $C_V(t,\tau) \geqslant 0$.

We also determine the prior loss function for informants, which will be useful later. Like villains, informant losses come from both victimization and punishment for failed crimes. The victimization losses of informants are identical to those of paladins, while the punishment losses are identical to those of villains. Hence, the strictly construed informant losses $\ell_I^-(t,\tau)$ are given by

$$\ell_I^-(t,\tau) = \frac{1}{\tau}\int_{t-\tau}^{t} (I+V)(A+V)(\delta+\epsilon) + (P+I)^2\theta \, ds \geqslant 0. \tag{3.10}$$

Rather than integrals, we can represent the loss or gain functions of equations (3.5), (3.6) and (3.9) as delay differential equations. Upon differentiation of the above equations and elimination of I, we find

$$\dot{\ell}_P(t,\tau) = \frac{\delta+\epsilon}{\tau}\{[1-P(s)-A(s)][A(s)+V(s)]\}\Big|_{t-\tau}^{t}, \tag{3.11}$$

$$\dot{\ell}_A(t,\tau) = \frac{\delta}{\tau}\{1-P(s)-A(s)\}\Big|_{t-\tau}^{t}, \tag{3.12}$$

$$\dot{\ell}_V(t,\tau) = \frac{\delta}{\tau}\{1-P(s)-3A(s)-2V(s)+[A(s)+V(s)]^2\}\Big|_{t-\tau}^{t} + \frac{\theta}{\tau}\{[1-A(s)-V(s)]^2\}\Big|_{t-\tau}^{t}. \tag{3.13}$$

Note that for the choice $\tau = 0$ in targeted recruitment, the loss functions greatly simplify, and are given directly by the integrands in equations (3.5), (3.6) and (3.9).

The objective functional that we seek to minimize is the total cost per citizen over the time horizon of interest $T > \tau$. As stated previously, this total cost includes both recruitment and carrying costs. The per capita recruitment costs $\mathcal{J}_r$ depend on which

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

recruitment strategy is used, and are given by

$$\mathscr{J}_r = \int_0^T dt \sum_{X\in\{P,A,V\}} m\, c\, (2-\ell_X) \begin{cases} u_X & \text{for targeted recruitment} \\ X u_X & \text{for untargeted recruitment.} \end{cases} \tag{3.14}$$

The per capita carrying costs $\mathscr{J}_c$ are found by combining the integrands of ℓ_P, ℓ_A, ℓ_V^- and ℓ_I^- with weights P, A, V and I respectively, then by integrating this quantity over the time horizon T. Hence, with I eliminated, the per capita carrying costs are

$$\mathscr{J}_c = \int_0^T (1-P-A)\left[(A+V)\delta + (1-A-V)(A+V)(\delta+\epsilon)\right]\, dt. \tag{3.15}$$

In the definition of $\mathscr{J}_c$, we have ignored the losses incurred by criminals due to punishment for failed crimes. If it was desired to minimize these punishment losses as well, an additional term

$$\theta(1-P-A)(1-A-V)^2$$

would be added to the integrand of $\mathscr{J}_c$. In addition, one could also consider carrying costs that may be associated with other aspects of criminal activity, such as costs of performing investigations, costs of punishing individuals (incurred by society, rather than by the person being punished) or others. For simplicity, however, we do not include any such terms here, but simply note that their inclusion could alter the results significantly.

Our optimal control problem, then, is to find the control variables $u_P(t)$, $u_A(t)$ and $u_V(t)$ that minimize the total cost $\mathscr{J} = \mathscr{J}_r + \mathscr{J}_c$ over the time horizon T:

$$\underset{u_P,u_A,u_V}{\operatorname{argmin}}\, [\mathscr{J}], \tag{3.16}$$

subject to the systems of equations (3.3) and (3.11)–(3.13), the control constraints in either equation (3.1) or (3.2) (depending on the type of recruitment employed) and the history functions

$$P(t) = 0,\ A(t) = A_d,\ V(t) = V_d \text{ for } t \leqslant 0, \tag{3.17}$$

so that the system begins in dystopia. The remainder of this paper will be devoted to this task and to the discussion of the solutions found.

4 Solutions

As our problem contains a large number of parameters, we choose to focus mainly on the effects of only two of them: c and m. All other parameter values are fixed for the remainder of this work, unless otherwise noted, and are given in Table 1. However, we will explore how changes to these parameters affect the optimal solutions in at least some limiting cases. The baseline values given in Table 1 are chosen to cover the widest range of optimal solutions in these limiting cases, as will be explored below. They are therefore not meant to necessarily represent 'realistic' values, but to allow us to explore the most interesting scenarios in our system.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

Table 1. *Parameter values used unless otherwise stated*

Parameter	Value	Description
δ	0.3	Payoff gain (loss) to the victimizer (victim).
ϵ	0.2	Additional victim loss due to retaliation.
θ	0.6	Severity of punishment.
τ	1	Memory length in untargeted recruitment.
T	100	Time horizon.

4.1 Infinite recruitment rate: $m \to \infty$

We begin with a limiting case, in which the recruitment rate m approaches infinity. Here all recruitment essentially happens instantaneously at $t = 0$ so that the system may then evolve towards utopia with no further external influence. In this case, only apathetics and villains may be recruited, since society contains no paladins to convert at $t = 0$. Since the recruitment in this case is instantaneous, the relevant 'control parameters' here are the total numbers of apathetics and villains recruited, denoted as U_A and U_V respectively. These quantities are expressed in dimensionless units so that U_X is really the total number of citizens of strategy type X converted to the informant role as a fraction of N, the total population.

In the infinitely fast recruitment case ($m \to \infty$), then, the optimal control problem is written as

$$\underset{U_A,U_V}{\operatorname{argmin}} \; [\mathcal{J}_{r,A} + \mathcal{J}_{r,V} + \mathcal{J}_c], \tag{4.1}$$

where the recruitment costs for apathetics, $\mathcal{J}_{r,A}$, and villains, $\mathcal{J}_{r,V}$, are functions of U_A and U_V that will be detailed below, and where $\mathcal{J}_c$ is found via equation (3.15) and the system of equation (3.3) with initial conditions $P(0) = 0$, $A(0) = A_d - U_A$ and $V(0) = V_d - U_V$.

4.1.1 *Targeted recruitment*

Consider first the case of targeted recruitment. Here the system exhibits no memory. Hence, even though recruitment happens at an essentially infinite rate, the costs to recruit individuals varies from one to the next. Because of this, it is important to consider the order of the recruitment; that is, if the optimal solution includes conversion of both apathetics and villains, is it better to recruit the apathetics or the villains first? Examining the loss functions above, we note that as villains are converted, villain losses increase, but the losses suffered by apathetics do not change, since the sum $I + V$ remains constant during the villain recruitment. On the other hand, as apathetics are converted, both villain losses and apathetic losses increase. Moreover, the losses to villains increase more rapidly when apathetics are converted than when villains are converted. Since increased losses lead to lower recruitment costs, and since the carrying costs are unaffected by the order of recruitment in this limiting case, it is clearly optimal to recruit apathetics first in this case (assuming they are recruited at all), as one may then reap the benefits of reduced conversion costs for both apathetics and villains.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

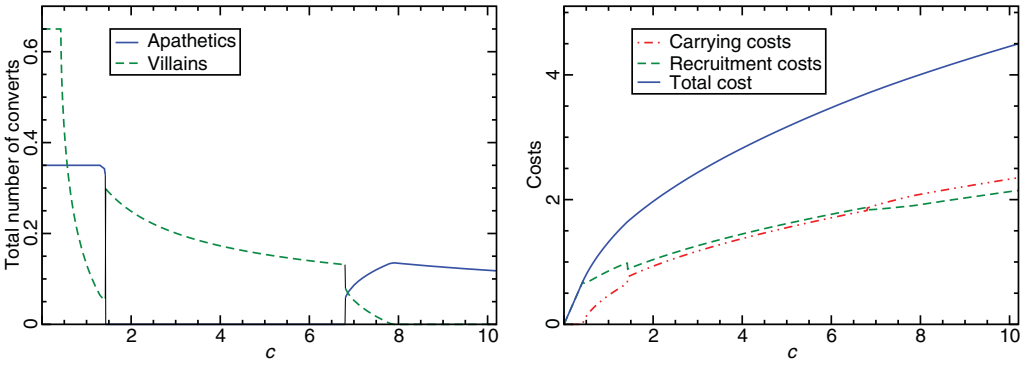


FIGURE 1. (Colour online) Results for our optimal control problem, equation (4.1), in the limiting case of an infinite recruitment rate $m \to \infty$, and using targeted recruitment. Left: The optimal number of apathetics (U_A) and villains (U_V) recruited at $t = 0$ as a function of the cost prefactor c. Note the discontinuous jumps in the solutions at $c \approx 1.4$ and $c \approx 6.8$. Right: The optimal carrying costs ($\mathscr{J}_c$), recruitment costs ($\mathscr{J}_r$) and total cost ($\mathscr{J}$) as functions of c.

With this in mind, we now calculate the targeted recruitment costs in terms of variables c, U_A and U_V. As mentioned above, if any apathetics at all will be converted to informants, they will be recruited first and at rate m, which is assumed very large. Hence, during apathetic recruitment, we have

$$P(t) = 0,\ A(t) = A_d - mt,\ I(t) = mt \text{ and } V(t) = V_d,$$

and the recruitment of U_A apathetics is completed over the time interval U_A/m. Upon substituting these functions for the state variables into the loss function $\ell_A(t, 0)$, which is given by the integrand of equation (3.6), we find via equation (3.14) that the recruitment cost for apathetics is

$$\mathscr{J}_{r,A} = cU_A\left[2 - \frac{\delta(1 + \delta + U_A)}{2}\right]. \tag{4.2}$$

We can similarly find the total recruitment costs for villains. In this case, since any apathetic recruitment that may occur happens before villains are recruited, we can write

$$P(t) = 0,\ A(t) = A_d - U_A,\ I(t) = U_A + mt \text{ and } V(t) = V_d - mt$$

during the recruitment of U_V villains, which occurs over the time interval U_V/m. Substituting the above expressions into $\ell_V(t, 0)$, given by the integrand of equation (3.9), we find

$$\mathscr{J}_{r,V} = cU_V\left[2 + \frac{\delta(1 - \delta - 2U_A)}{2} - (\delta + \theta)\left(U_A^2 + U_A U_V + \frac{U_V^2}{3}\right)\right]. \tag{4.3}$$

With the targeted recruitment cost functions identified, we can proceed to solve our control problem, equation (4.1). For a given value of c, this problem can be solved using various techniques; we chose the *Mathematica* function 'NMinimize'. Plots of the number of converted individuals U_A and U_V and of the optimal costs $\mathscr{J}_c$, $\mathscr{J}_r$ and $\mathscr{J}$ as functions of c in the case of targeted recruitment are shown in Figure 1.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

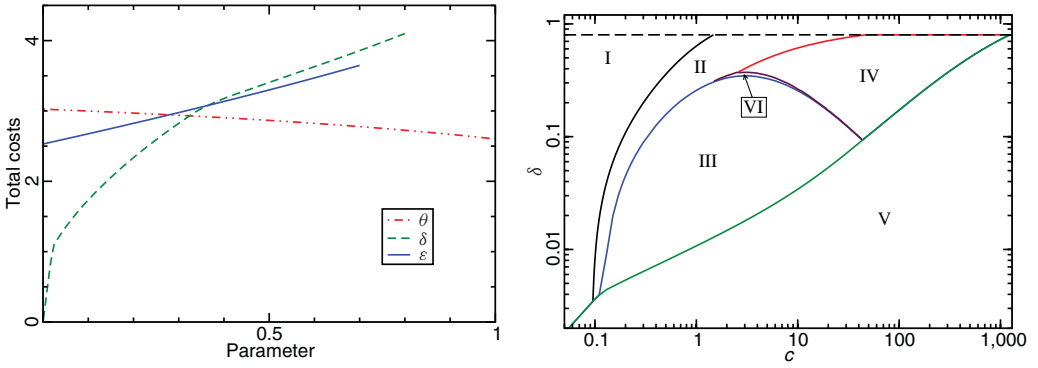


FIGURE 2. (Colour online) Sensitivity analysis of the optimal solution as parameters θ, δ and ϵ are varied from the values given in Table 1 using targeted recruitment. Left: The optimal total costs as one parameter is varied at a time, using the values from Table 1 for the other two parameters and $c = 4$. As θ and ϵ vary, the qualitative optimal strategy – convert only villains – does not change, but as δ varies, the optimal strategy takes on five qualitatively distinct forms. Right: A phase diagram in (c, δ) space, displaying the regions where the six qualitatively distinct optimal solutions are obtained; these solutions I–VI are described in the text.

Referring to Figure 1, we find several behaviours that will be discussed in detail in Section 4.1.3; for now, we simply make note of them. First, as $c \to 0$, all citizens are recruited. Second, as c crosses ≈ 0.3, we begin recruiting fewer villains, but not fewer apathetics. Third, we see that the number of individuals recruited of each type behaves discontinuously at two points in the c domain. At $c \approx 1.4$, the behaviour switches from that of converting all available apathetics and some villains to converting only villains. We see that in the optimal costs, a tradeoff has been made here between reducing recruitment costs and increasing carrying costs. At $c \approx 6.8$ the behaviour switches back to converting some members of each category. Here we see that the recruitment and carrying costs curves have crossed, with recruitment now becoming the lesser of the two. Fourth, we see that as c becomes large ($c \gtrsim 7.8$), we recruit fewer and fewer citizens, all of which are apathetics. Finally, though it is not shown in Figure 1, for $c \gtrsim 220$ the optimal solution is to recruit no one.

At this point we might ask how these solutions change with different choices of the parameters θ, δ and ϵ. The results of such a sensitivity analysis are presented in Figure 2. The left half of Figure 2 shows how the total optimal costs $\mathcal{J}$ vary as we change the value of one parameter at a time, keeping the others fixed at their values from Table 1 and using $c = 4$. We observe that the parameters $0 \leqslant \theta \leqslant 1$ and $0 \leqslant \epsilon \leqslant 1 - \delta$ have relatively minor effects on the optimal costs, with increasing θ causing lower costs and increasing ϵ causing increased costs. These two effects are clear: Increasing θ indirectly decreases the carrying costs (as punished criminals convert to non-offenders more quickly) and directly decreases the recruitment costs of villains; increasing ϵ both directly and indirectly (as failed witnesses more quickly convert to criminals) increases the carrying costs, and has no explicit connection to the recruitment costs. Importantly though, varying either of these parameters separately does not qualitatively change the optimal solution for this c value, to convert only villains, as seen in Figure 1.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

In contrast, changing $0 \leqslant \delta \leqslant 1 - \epsilon$ has a relatively large effect on the optimal costs, with increasing δ leading to increased costs. This effect is understandable, as increasing δ both directly and indirectly (as victims more quickly convert to criminals) increases the carrying costs, thereby necessitating the conversion of more citizens, though at a generally lower cost per convert. Furthermore, as δ varies from its lowest to the highest possible values, the qualitative optimal strategy takes on five different forms, leading to dramatically different optimal solutions. This effect is illustrated on the right side of Figure 2, where the regions in which six qualitatively different optimal solutions obtained are shown in (c, δ) space, with all other parameters given in Table 1. Here the six different solution types are as follows:

- I Convert all citizens.
- II Convert all apathetics and some, but not all, villains.
- III Convert no apathetics and some, but not all, villains.
- IV Convert some, but not all, apathetics and no villains.
- V Convert no one.
- VI Convert some, but not all, apathetics and villains.

Hence, for $c = 4$, as δ varies from its lowest to the highest values the qualitative solution takes on the forms V, III, VI, IV and II, in that order. Similarly, using the value of δ (0.3) from Table 1 and varying c brings the system through the solutions I, II, III, VI, IV and V, in that order, exactly as seen in Figure 1. Note that solutions I–V exist on the boundary of the permissible region of (U_A, U_V) space, while only solution VI is an interior minimum of the total costs. We shall discuss these solutions, and the boundaries between them, more in Section 4.1.3.

4.1.2 *Untargeted recruitment*

Now consider the case of untargeted recruitment, in which the system exhibits memory. Since $\tau > 0$ but recruitment is happening instantaneously, each apathetic converted into an informant bears the same recruitment cost, which is solely dictated by the fact that the system has been in dystopia, and therefore at equilibrium, for all times $t \leqslant 0$. The same statement holds true for villains as well. Thus, in this case recruitment costs for apathetics and villains, found through equation (3.14), are simply given by

$$\mathscr{J}_{r,A} = cU_A \left[2 - \frac{\delta(1+\delta)}{2}\right], \quad \mathscr{J}_{r,V} = cU_V \left[2 + \frac{\delta(1-\delta)}{2}\right]. \tag{4.4}$$

With the recruiting costs found, we now solve the problem as in the targeted recruitment case; results are displayed in Figure 3. We see that the behaviour of the optimal solution here is qualitatively very similar to that of the targeted recruitment. The only qualitative difference occurs at the second discontinuity in recruitment values, at $c \approx 8.8$. In the untargeted case, this discontinuity causes a shift from converting only villains (solution type III) to converting only apathetics (solution type IV), while in the targeted case the corresponding discontinuity causes a shift from converting only villains (solution type III) to converting both villains and apathetics (solution type VI).

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

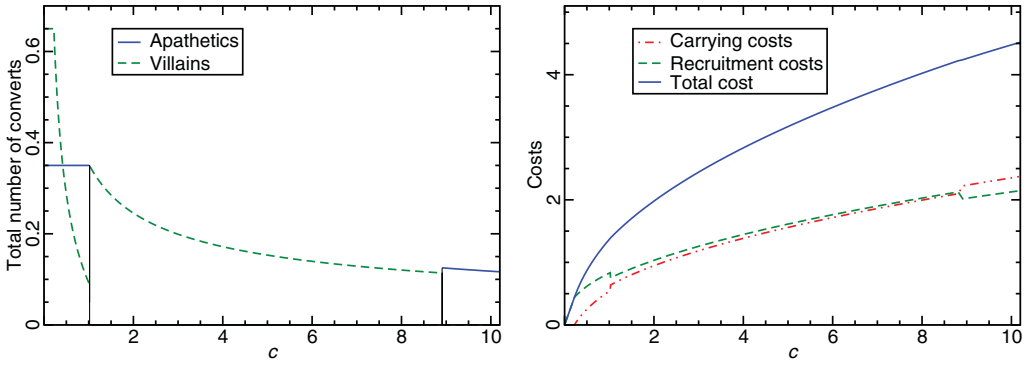


FIGURE 3. (Colour online) Results for our optimal control problem, equation (4.1), in the limiting case of an infinite recruitment rate $m \to \infty$, and using untargeted recruitment. Left: The optimal number of apathetics (U_A) and villains (U_V) recruited at $t = 0$ as a function of the cost prefactor c. Note the discontinuous jumps in the solutions at $c \approx 1.0$ and $c \approx 8.8$. Right: The optimal carrying costs ($\mathscr{J}_c$), recruitment costs ($\mathscr{J}_r$) and total cost ($\mathscr{J}$) as functions of c.

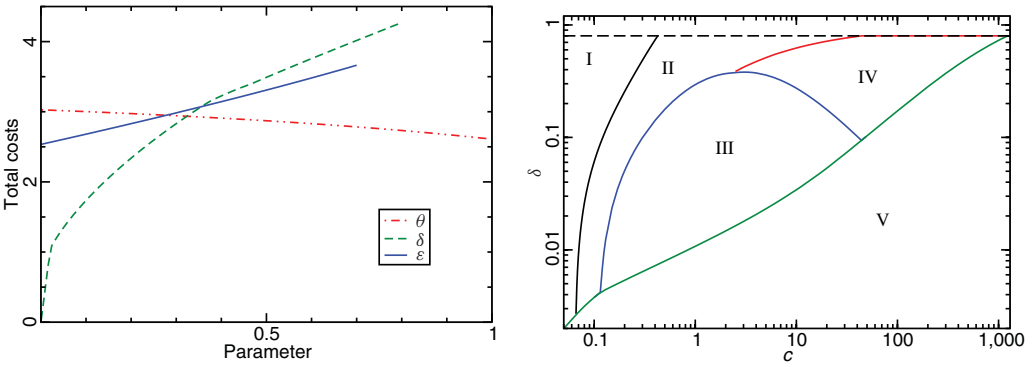


FIGURE 4. (Colour online) Sensitivity analysis of the optimal solution as parameters θ, δ and ϵ are varied from the values given in Table 1, using untargeted recruitment. Left: The optimal total costs as one parameter is varied at a time, using the values from Table 1 for the other two parameters and $c = 4$. As θ and ϵ vary, the qualitative optimal strategy – convert only villains – does not change, but as δ varies, the optimal strategy takes on four qualitatively distinct forms. Right: A phase diagram in (c, δ) space, displaying the regions where five qualitatively distinct optimal solutions are obtained; these solutions I–V are described in the text.

This difference in behaviour between the targeted and untargeted case can also be found in the results of a sensitivity analysis performed on the untargeted case, similar to that performed for the targeted case above (Figure 2). Plots for this untargeted sensitivity analysis are shown in Figure 4. The results in this case are overall very similar to those seen for the targeted case, with one major exception – no solutions of type VI exist in the case of untargeted recruitment. Hence, as δ varies from zero to $1 - \epsilon$ on the left side of Figure 4, the optimal strategy takes on only four values: V, III, IV and II, in that order.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

And, for $\delta = 0.3$ and varying c, the untargeted system goes through solution regions I, II, III, IV and V, as seen in Figure 3.

Despite the qualitative differences between targeted and untargeted recruitment, we find that the optimal costs for any given combination of parameters in each case to be quite similar quantitatively. If we compare the costs illustrated in Figures 1 and 3, for example we find that the total costs differ by less than a few percent over most c values. The exception to this is at very low c, where the solution in both cases is of type I and the differences in costs can rise to around 20%. However, even in these cases, the absolute difference in optimal costs between the two recruitment types is less than 0.1 – approximately one-third of the losses from one successful victimization.

4.1.3 *Discussion*

In this section, we explain the many behaviours observed in Figures 1–4 to gain intuition about how our optimal strategies are being selected. We begin by examining two extreme cases: solution type V (convert no one) and solution type I (convert everyone). For type V, only carrying costs exist, and since the system does not evolve from dystopia in this case, the total costs are

$$\mathscr{J} \equiv \mathscr{J}_{\mathrm{V}} = \frac{\delta(1+\delta)T}{2}. \tag{4.5}$$

Alternatively, for solution type I, only recruitment costs exist. Thus, if targeted recruitment is used, we have a total cost of

$$\mathscr{J} \equiv \mathscr{J}_{\mathrm{I}} = \frac{c}{24}[48 - 7\theta - \delta(16 - 3\delta - 6\delta^2 + \delta^3) - \delta\theta(3 - 3\delta + \delta^2)], \tag{4.6}$$

while if untargeted recruitment is used we have the significantly simpler

$$\mathscr{J} \equiv \mathscr{J}_{\mathrm{I}} = 2c. \tag{4.7}$$

Hence, for a given value of $\delta > 0$, we clearly see that, as $c \to 0$, the optimal solution will be of type I, as that choice will lead to $\mathscr{J} \to 0$, the smallest it could possibly be. Alternatively, for a fixed value of $c > 0$, as $\delta \to 0$, the optimal solution will clearly be of type V, as that choice will lead to $\mathscr{J} \to 0$. Furthermore, the boundary between these two extreme solutions is a portion of the line in (c, δ) space along which the costs in equation (4.5) are equal to those in equation (4.6) or (4.7), depending on the recruitment method used; below this curve, type V is a better solution, and above it, type I is better (though another solution type might be better than either). This argument explains the behaviours on the very left and on the very bottom of the phase plots in Figures 2 and 4, and the boundary between regions I and V on these same plots.

Now consider what happens when we deviate slightly from a type V solution, and begin converting a very small number of citizens to informants. In this regime, the approximation in equation (2.7) for the time to utopia, T_u, will hold. During most of this time, the rate at which carrying costs are accrued is actually higher than if the system had been left in the pure dystopian state. This is because we have introduced a small number of citizens who will report when victimized, but who will almost never be successful in doing so, and thereby suffer the addition loss ϵ compared with their non-reporting counterparts.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

Furthermore, if the small number of converts are apathetics rather than villains, there are a larger total number of criminals in the system, leading to an even higher increased rate in this case. Because of these facts, converting a small number of citizens is only beneficial if the system reaches Utopia, where carrying costs are accrued at a rate of zero before the time horizon is reached. Hence, for initial informant values $I_0 \lesssim 2/T(2-\epsilon)$ (for which $T_u \gtrsim T$) carrying costs are actually increasing in U_A and U_V. In fact, one can show that the marginal carrying costs at the type V solution state are given by the positive functions

$$\frac{\partial \mathscr{J}_c}{\partial U_V} = \frac{T\epsilon(2+\delta-\epsilon)}{2} - \frac{\epsilon(1-\epsilon)(2-\epsilon)}{(1+\delta)^2}\left[1 - e^{T(1+\delta)^2/2(\epsilon-2)}\right], \tag{4.8}$$

$$\frac{\partial \mathscr{J}_c}{\partial U_A} = \frac{\partial \mathscr{J}_c}{\partial U_V} + \frac{2\delta}{1+\delta}\left[1 - e^{-T(1+\delta)/2}\right]. \tag{4.9}$$

Since the marginal recruitment costs are also positive at the type V solution (regardless of recruitment type), the type V solution is always a local minimum. In fact, as $c \to \infty$, solution type V becomes the global minimum, as its cost does not change, but the costs of all other possible solutions grow very large. This explains the behaviour on the far right of the phase plots in Figures 2 and 4.

Now consider a fixed $\delta > 0$ and $c \gg 1$. Since c is large, conversion costs dominate, and the optimal solution will be to either convert only a very small number of citizens or none at all (type V). If we assume for now that the local minimum at solution type V is not the global minimum (c is not too large), then we can solve for the optimal solution in the following way. As stated above, the system in this case will spend an approximate amount of time T_u very near the dystopian state. Therefore, the carrying costs can be approximated through equation (4.5), but substituting T_u for T, since we are assuming the system reaches Utopia before the time horizon; this gives

$$\mathscr{J}_c \approx \frac{\delta(1+\delta)}{(2-\epsilon)(U_A+U_V)}. \tag{4.10}$$

According to this formula, the carrying costs do not depend on which types of citizens were recruited, but only the total number of recruits. But, as explained above, the carrying costs are slightly higher when apathetics are converted instead of villains, mainly due to the increased number of criminals. Though some of this increased cost is seen in the extra term appearing in equation (4.9), that extra factor is quite small, as it is proportional to U_A. Numerical studies show that a similar, but much larger, increase arises during the brief period over which the system transforms from the near dystopian state to the utopian state, and that this increase is approximately (for the ϵ and θ values chosen) $11\delta/10$ times the fraction of initial recruits that were apathetics; it therefore does not directly depend on the total number of recruits. Hence, using untargeted recruitment, the total costs in this case will be

$$\begin{aligned}\mathscr{J} \approx{}& \frac{\delta(1+\delta)}{(2-\epsilon)(U_A+U_V)} + cU_A\left[2 - \frac{\delta(1+\delta)}{2}\right] \\ &+ cU_V\left[2 + \frac{\delta(1-\delta)}{2}\right] + \frac{11\delta U_A}{10(U_A+U_V)}.\end{aligned} \tag{4.11}$$

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

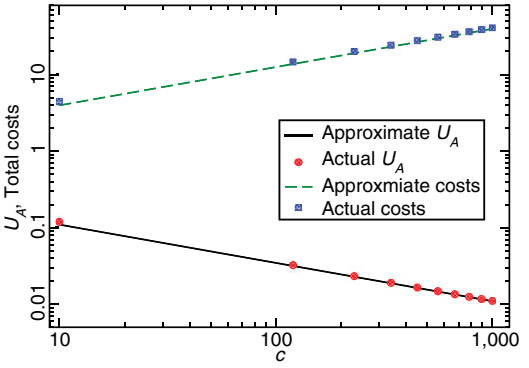


FIGURE 5. (Colour online) Plots of the numerically determined optimal U_A and $\mathscr{J}$, as well as the approximate values given in equation (4.12), versus large c values in the $m \to \infty$ case, with $T = 10,000$. Note that this is on a log–log scale. The actual U_A is well approximated by the power law $U_A \sim 1/\sqrt{c}$, while the actual total cost $\mathscr{J}$ is well approximated by the power law $\mathscr{J} \sim \sqrt{c}$.

Using these approximated costs, we can minimize over U_A and U_V, finding that no interior minimums exist, and that if apathetics are converted, the minimum is

$$U_A \approx \sqrt{\frac{2\delta(1+\delta)}{c(2-\epsilon)\,[4-\delta(1+\delta)]}}\,, \quad \mathscr{J} \approx \sqrt{\frac{2c\delta(1+\delta)\,[4-\delta(1+\delta)]}{2-\epsilon}} + \frac{11\delta}{10}\,, \tag{4.12}$$

while if villains are converted, the minimum is

$$U_V \approx \sqrt{\frac{2\delta(1+\delta)}{c(2-\epsilon)\,[4+\delta(1-\delta)]}}\,, \quad \mathscr{J} \approx \sqrt{\frac{2c\delta(1+\delta)\,[4+\delta(1-\delta)]}{2-\epsilon}}. \tag{4.13}$$

Hence, we find that, for c large, our solution is $U_X \sim 1/\sqrt{c}$ and $\mathscr{J} \sim \sqrt{c}$; Figure 5 shows the legitimacy of this approximate scaling for large c values, using $T = 10,000$ to avoid solution type V being a global minimum.

Importantly, though, the costs above allow us to determine whether our system will prefer to convert only a small number of villains (type III), a small number of apathetics (type IV) or no one (type V), at least at large values of c. The boundary between types III and IV solutions at high c occurs when the costs in equations (4.12) and (4.13) are equal. Below this line in the (c, δ) space, the greater recruitment costs of villains are made up for by their lesser carrying costs, making them the optimal choice; above this line, the increased carrying costs of apathetics are made up for by their lesser recruitment costs, so they are the optimal choice. The boundaries between either of these solutions and type V are found by equating the cost from equations (4.12) or (4.13) (whichever is lesser) to the cost in equation (4.5). The approximate boundaries derived by these ways are illustrated in Figure 6, and compare favourably to those seen in the phase plot of Figure 4, thus explaining these portions of our solution space.

This argument also holds for targeted recruitment, but due to the considerably more complicated forms of the targeted recruitment costs, analytically determining the boundaries between regions is unfeasible. However, we do find in this case that, because of the

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

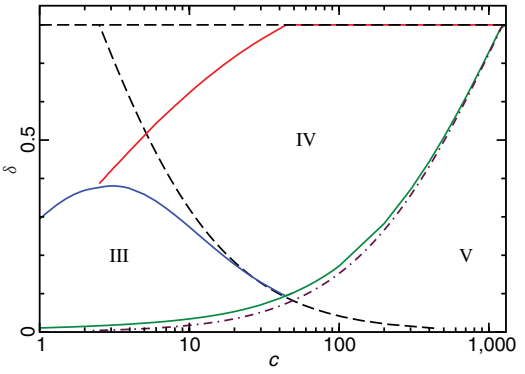


FIGURE 6. (Colour online) Comparing the approximate boundaries (dashed lines) between regions III, IV, and V in untargeted recruitment to the actual boundaries (solid lines) from Figure 4. The approximate boundaries compare favourably to the actual curves for large c values.

nonlinear recruitment costs, our approximate total cost admits minimums that are interior solutions of type VI, explaining their existence in the case of targeted recruitment. These solutions exist in a narrow band in (c, δ) space surrounding the curve along which the types IV minimum and III minimum are equal in costs, explaining why region VI is so small in our phase plot.

At this point we turn to the solution type II, the only type whose existence has not yet been explained. Let us consider a fixed $\delta > 0$ and $c \ll 1$. Furthermore, assume we are above the curve separating regions I and V, explained above, so that the type V solution is certainly not the global minimum. Here, since c is quite small, carrying costs dominate so that the possible solutions are to convert all citizens (type I) or almost all citizens. If the latter solution is the global minimum, it must be the case that we realize a net savings by discontinuing the conversion of a small number of citizens from the type I solution, thus increasing carrying costs but decreasing conversion costs. In untargeted recruitment, more is saved on recruitment costs by discontinuing the conversion of villains, rather than apathetics, with the marginal recruitment savings for villains being

$$\frac{\partial \mathscr{J}_r}{\partial U_V} = c\left[2 + \frac{\delta(1-\delta)}{2}\right]. \tag{4.14}$$

Interestingly, in this setting there is also a smaller increase in carrying costs by discontinuing the conversion of villains, rather than apathetics. This is because the system begins in this case with almost all citizens willing to serve as witnesses, hence almost all crimes are unsuccessful, except for those against the few unconverted citizens, who do not report the crime in the first place. But, since villains are more active than apathetics – they can be both victims and criminals, in which case they will almost certainly be punished – those few villains left behind convert more quickly to either an informant or paladin (against whom a crime will almost certainly be unsuccessful) than would similarly left behind apathetics, leading to lower comparative carrying costs. In fact, the marginal carrying cost increases arising from the non-conversion of apathetics or villains from the type I

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

solution can be calculated exactly, and are found to be

$$-\frac{\partial \mathscr{J}_c}{\partial U_A} = \frac{2(2\delta+\epsilon)}{1+\delta}\left[1-e^{(1+\delta)(\theta-2)/2}\right], \tag{4.15}$$

$$-\frac{\partial \mathscr{J}_c}{\partial U_V} = \frac{2(2\delta+\epsilon)\left[(1-\delta)(1-\theta)+e^{(1+\delta)(\theta-2)/2}\left[\delta(2-\theta)+\theta\right]-e^{\theta-2}(1+\delta)\right]}{(1-\delta^2)(2-\theta)}. \tag{4.16}$$

Therefore, once c is large enough that the marginal recruitment savings for villains in equation (4.14) exceeds the marginal carrying cost increases for villain non-conversion found in equation (4.16), the optimal solution is to recruit all apathetics and some, but not all, villains, thus explaining (and giving the functional form of) the line separating regions I and II on the phase diagram in Figure 4.

The same argument also holds true for targeted recruitment, with the modification that the marginal recruitment savings for villains in solution type I,

$$\frac{\partial \mathscr{J}_r}{\partial U_V} = c(2-\delta-\theta), \tag{4.17}$$

are no longer always greater than those for apathetics,

$$\frac{\partial \mathscr{J}_r}{\partial U_A} = \frac{c}{4}\left[8-\delta(9+4\delta-\delta^2)-\theta(3-\delta)(1+\delta)\right], \tag{4.18}$$

especially at small values of δ. However, despite this, the overall savings are still always greater by discontinuing the conversion of villains rather than apathetics (due to the lower increase in carrying costs for villains) so that at the critical c value, now found by equating equation (4.17) with equation (4.16), we will again stop recruiting all villains, but still recruit all apathetics.

4.2 Finite recruitment rate

4.2.1 *Targeted recruitment*

In order to solve our system in the case of targeted recruitment, we made use of the freely available software package PSOPT [6], which uses direct pseudospectral methods to solve optimal control problems. Results showing total recruits U_X and the costs for various values of c and m are shown in Figure 7, while results for the optimal $u_X(t)$ for various values of c and m are shown in Figure 8.

The results in the case of targeted recruitment exhibit some behaviour that is expected, but also some behaviour that is somewhat surprising. One noteworthy finding is that, for $c \gtrsim 3.25$, paladins are always converted at low enough m values, highlighting the fact that informants are even more valuable in reforming a society than their law-abiding counterparts, the paladins. More interestingly, conversion of paladins seems to be favoured even over that of apathetics in such cases, as seen in the top left of Figure 8, where resources are being diverted from the conversion of apathetics towards that of paladins at early times t. Here the control u_P is less than one but increasing; since only a small number of paladins exist at early times, they must be recruited at a rate less than m lest P become negative. This favouring of paladins over apathetics is likely due to the

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

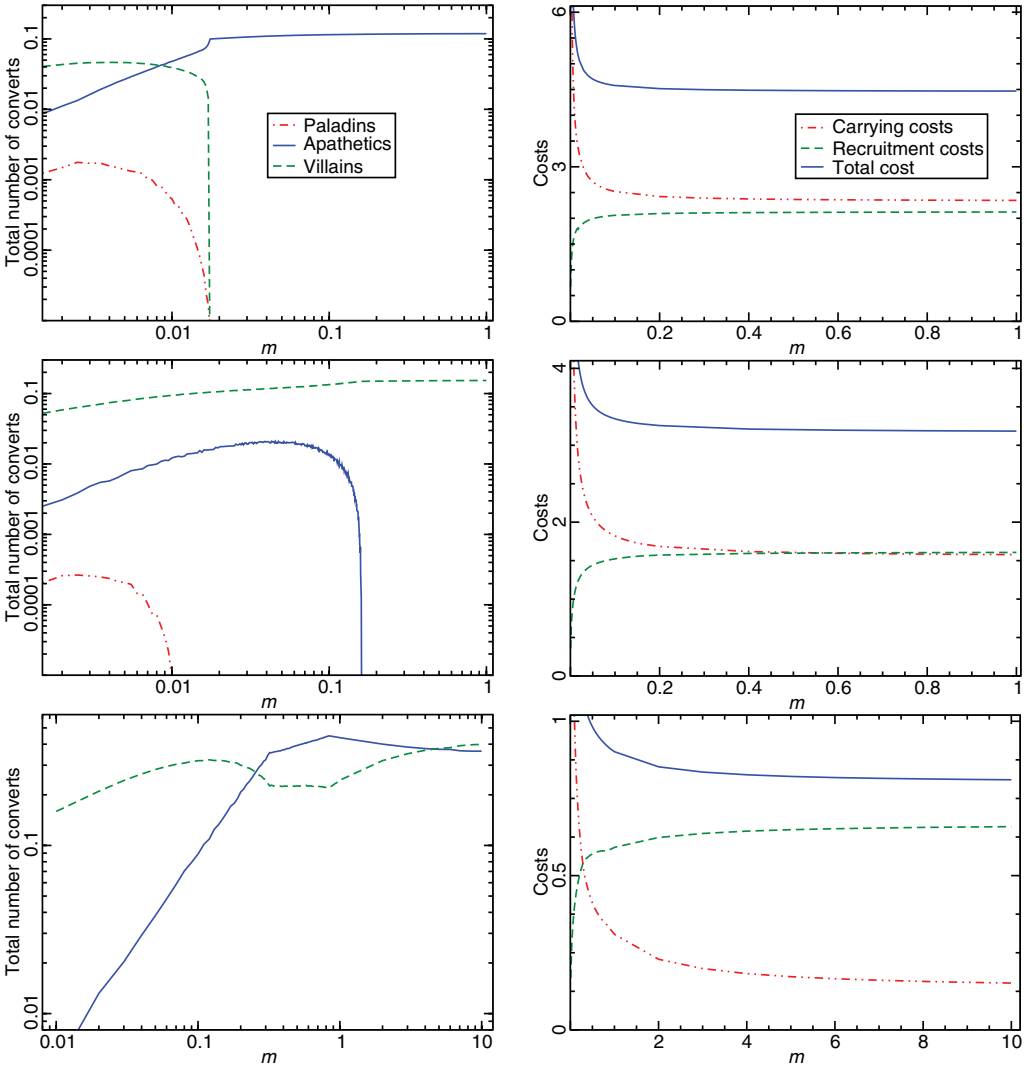


FIGURE 7. (Colour online) Results for our optimal control problem, equation (3.16), with targeted recruitment. The left column displays plots of the optimal total number of paladins (U_P), apathetics (U_A) and villains (U_V) recruited as functions of m for the three different choices $c = 10.0$ (top row), $c = 5.0$ (middle row) and $c = 0.5$ (bottom row). Note that these plots are on log–log scales. The right column displays plots of the optimal carrying costs ($\mathscr{J}_c$), recruitment costs ($\mathscr{J}_r$) and total cost ($\mathscr{J}$) for these same c values. For each c value, the solutions here approach those of Figure 1 as m becomes large.

fact that paladins are less costly to convert than apathetics near $t = 0$, when paladins are suffering both victimization and retaliation losses (apathetics only suffer victimization losses). This conversion cost difference therefore offsets any advantage that apathetics may have over paladins in terms of slightly reduced carrying costs. Of course, if m is large enough, there are enough resources available to convert many apathetics and villains so that paladins need not be converted any longer (see Figure 7, top and middle rows). For

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

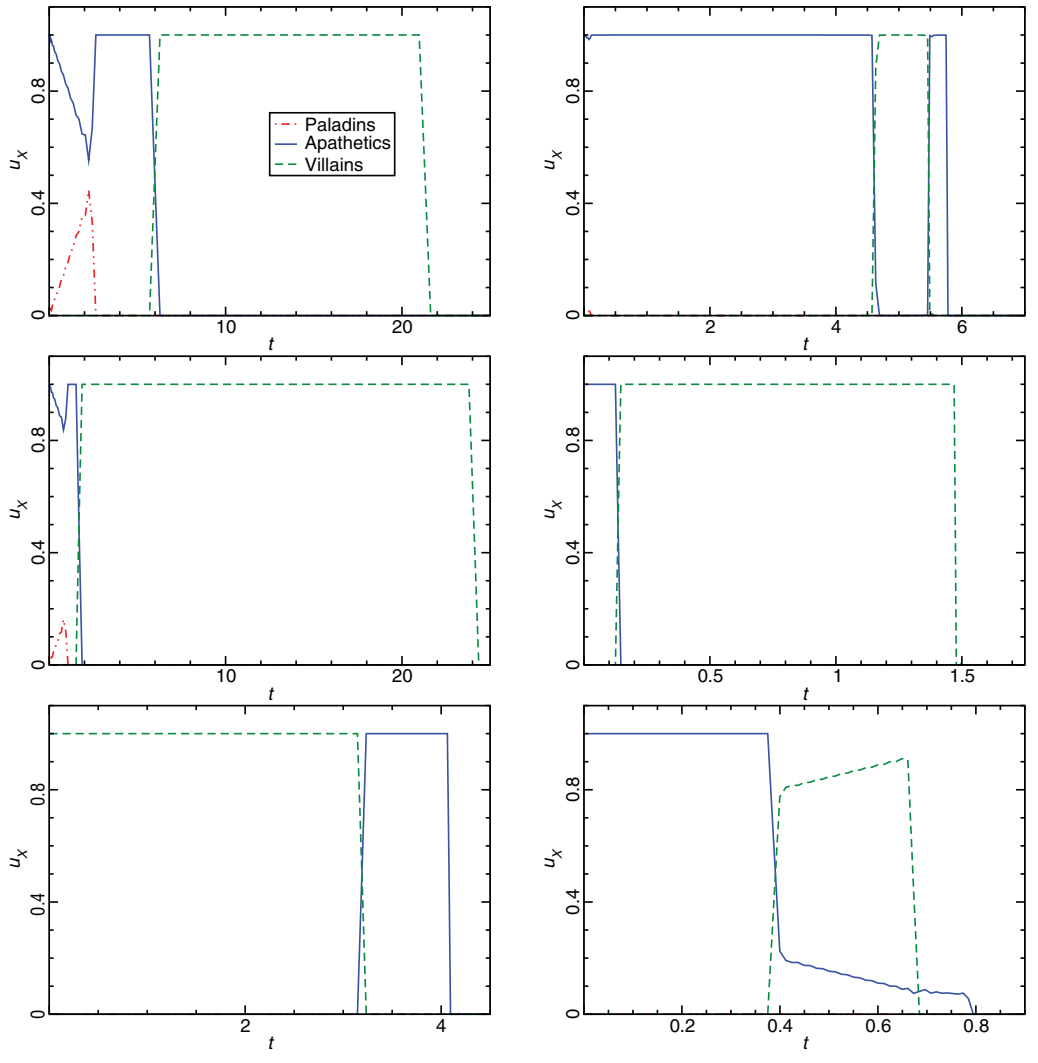


FIGURE 8. (Colour online) Plots of selected optimal controls $u_X(t)$ for our problem with targeted recruitment. Top row: Results with $c = 10.0$, using $m = 0.003$ (left) and $m = 0.017$ (right). Middle row: Results with $c = 5.0$, using $m = 0.003$ (left) and $m = 0.1$ (right). Bottom row: Results with $c = 0.5$, using $m = 0.1$ (left) and $m = 1.0$ (right).

$c \lesssim 3.25$, however, no paladins ever seem to be converted, regardless of m (see Figure 7, bottom row). This is presumably because the difference in costs between paladins and apathetics (or villains) has decreased enough in this regime to negate the advantage they held at higher c values.

Another finding of interest is that apathetic recruitment does not always come before villain recruitment in the case of small m, as opposed to the result when $m \to \infty$. For $c \gtrsim 3.25$, the same regime where paladins may also be converted, most apathetic conversion generally takes place before any villains are recruited, in cases where both types are indeed converted at all. But for smaller c and m values, as in the bottom left

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

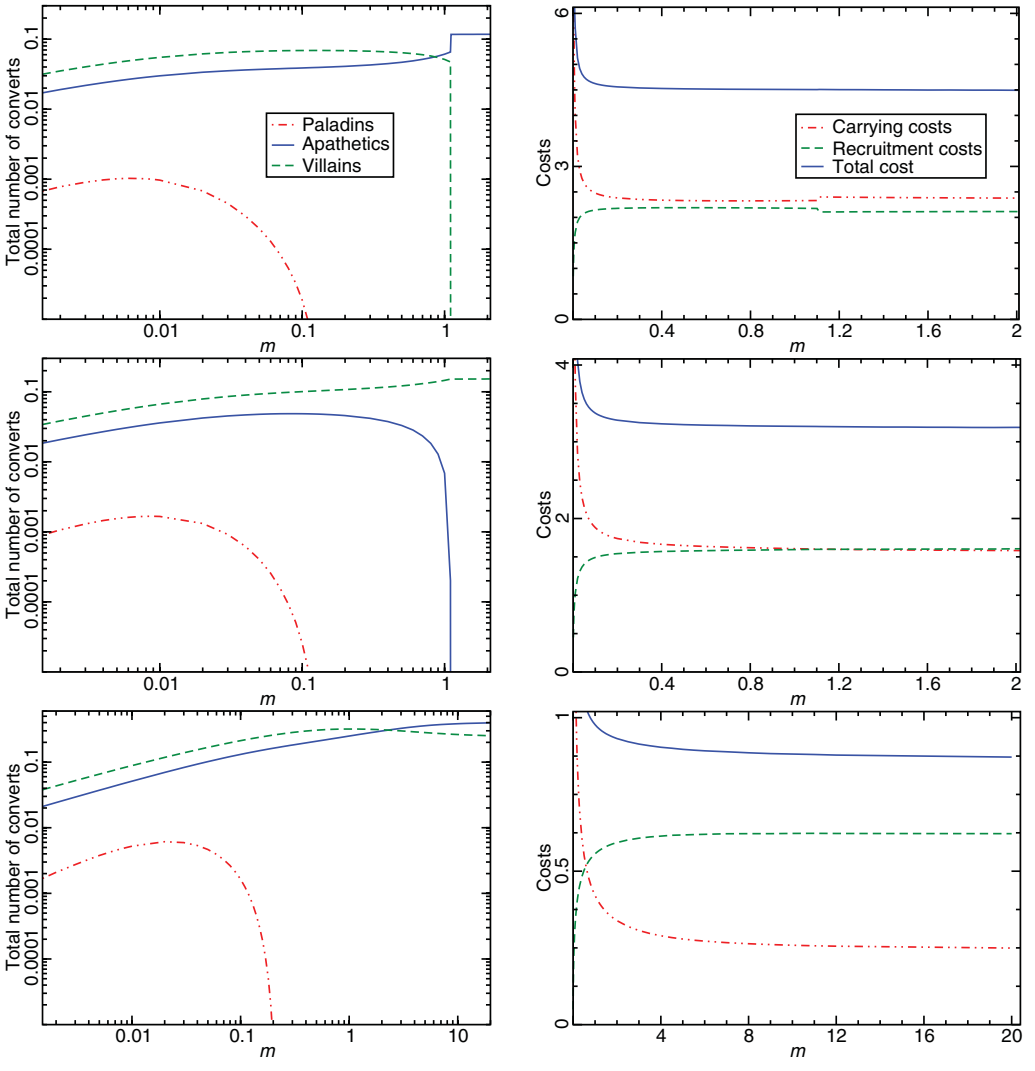


FIGURE 9. (Colour online) Results for our problem with untargeted recruitment. The left column displays plots of the optimal total number of paladins (U_P), apathetics (U_A) and villains (U_V) recruited as functions of m for the three different choices $c = 10.0$ (top row), $c = 5.0$ (middle row) and $c = 0.5$ (bottom row). Note that these plots are on log–log scales. The right column displays plots of the optimal carrying costs ($\mathscr{J}_c$), recruitment costs ($\mathscr{J}_r$) and total cost ($\mathscr{J}$) for these same c values. For each c value, the solutions correctly approach those of the $m \to \infty$ case as m becomes large.

of Figure 8, we see that villains are recruited before apathetics. We theorize that this occurs because of the advantage in reduced carrying costs for recruiting villains over apathetics discussed earlier, since recruiting villains does not increase the total number of criminals in the society. This advantage must therefore be greater than any recruitment cost decreases that may be obtained by converting apathetics first in this regime.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

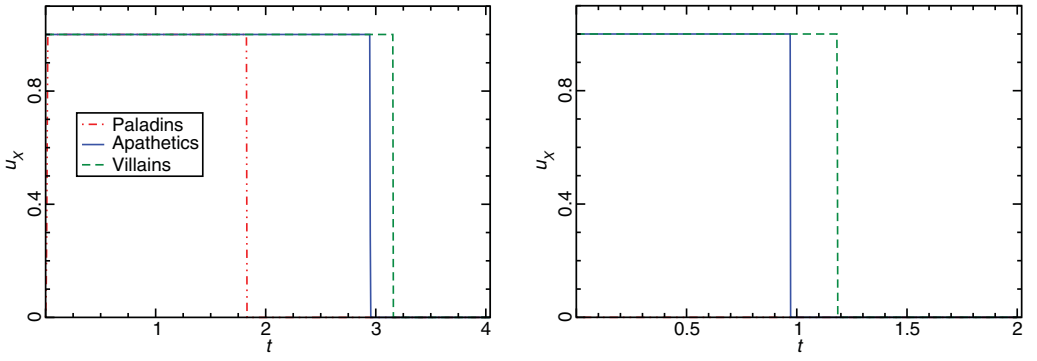


FIGURE 10. (Colour online) Plots of selected optimal controls $u_X(t)$ for our problem with untargeted recruitment and $c = 5.0$, using $m = 0.05$ (left) and $m = 0.15$ (right).

Finally, we note that the solutions for $u_X(t)$ occasionally exhibit a switching back and forth of the preferred convert type. We see such behaviour in the top right of Figure 8, where the solution is to convert only apathetics until $t \approx 4.6$, then to switch to only villains until $t \approx 5.5$, then to switch back to apathetics until $t \approx 5.75$. The take-home message here is that, under targeted recruitment, there are many subtle effects that can lead to non-trivial optimal solutions.

4.2.2 *Untargeted recruitment*

Due to technical difficulties involving the introduction of the delay differential equations, the solutions for the case of untargeted recruitment were found not with PSOPT, but by creating our own solver based on satisfying the necessary conditions for optimality (see Appendix A for a discussion of these conditions and a description of our numeric algorithm). Results showing total recruits U_X and the costs for various values of c and m are shown in Figure 9, while results for the optimal $u_X(t)$ for various values of c and m are shown in Figure 10.

The results for untargeted recruitment generally exhibit less structure than those in targeted recruitment. This is because in this case all three strategy types are potentially being recruited at any given time, whereas in targeted recruitment, favoured recruit types could be chosen and solely targeted. However, some interesting findings still arise in this case. For example, the solutions at small m are quite similar across different c values; contrast the three plots on the left of Figure 9 to observe this. Specifically, we find that all three strategy types are converted initially at low m. This is again because of the inability of authorities to only target specific citizen strategies for recruitment so that at low m the authorities cannot afford to be too picky about which strategies they choose to recruit, and will simply recruit everyone they can in order to quickly increase the number of informants. Because of this, the state variables evolve in a very regular way regardless of c, leading to the quite similar solutions at low m.

Comparing the costs of untargeted recruitment (Figure 9, right column) with those of targeted recruitment (Figure 7, right column) for the same c and m, we find only small differences, as noted in the $m \to \infty$ case above. This again seems to indicate that having

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

complete knowledge of current citizen strategies does not offer much of an advantage to the authorities in terms of societal reformation, and that the random recruitment strategy is nearly as good. Hence, our work shows that if the gaining of such detailed knowledge comes with any non-trivial price, it is likely not worth the cost.

5 Conclusion

In this paper we developed an optimal control problem to determine the most cost-effective way of recruiting informants from a general population with diverse attitudes towards committing and reporting crimes, with the goal of driving a criminal society, only populated by villains and apathetics, towards a crime-free, utopian state. We considered targeted and untargeted recruitment of citizens, where authorities may or may not have knowledge of individual strategy types, and analyzed the consequences of both assumptions. For both recruitment strategies, we introduced a total societal cost of crime and crime abatement as given by a combination of direct recruitment costs and losses to society due to criminal activity. Recruitment costs vary from individual to individual and – in the case of untargeted recruitment – on the potential recruit's payoff history. We find that, depending on the resources m at hand and on the magnitude of recruitment costs c, it may be more beneficial to focus the conversion efforts on some particular citizen types rather than on others, and according to a select ordering. We find that the most interesting results emerge in the case of targeted recruitment. Here authorities know in advance what types of citizens are being converted and this knowledge results in very nuanced behaviours. In particular, two interesting limiting cases emerge: that of a plentiful society where resources are abundant and recruitment costs are low (large m and small c), or the opposite case of a strained society where the resource pool is limited and conversions are costly (low m and large c).

In the first case (large m and small c), we were able to show that it is always preferable to recruit apathetic citizens first, since this tactic will lead to lower future costs in recruiting other apathetics or villains. Within this scenario, paladins are created as a consequence of the evolving dynamics but are never actively converted to informants. Close inspection of other cases shows that paladin recruitment is also not viable for values of m and c that are either both large or both small. However, the nonlinear interplay between costs and benefits does favour the conversion of paladins into informants in the case of a strained society (low m and large c). Here conversion of all citizen types occurs, with paladins being preferred over apathetics and villains, in this order. Less structured results are seen in the case of untargeted recruitment due to a less selective recruitment process that tends to equalize effects, regardless of costs or resources.

Upon comparing results from the targeted and untargeted cases, it is apparent that the outcome differences between the two recruitment methods are minor. Since in practice the implementation of targeted recruitment may be more challenging, due to related issues, such as privacy or profiling concerns, untargeted recruitment may be just as efficient a tool as targeted recruitment.

Finally, this work may be generalized in many ways. For instance, rather than attempting to purely minimize costs as we have done, an alternative goal could be to achieve some specified low-level crime state (not necessarily utopia) at the time horizon while using a

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

finite amount of resources. This is related to what we have done, but places the emphasis on the final state that is achieved. Similarly, one could consider different weights of the carrying and recruitment costs, depending on the particular society at hand. Spatial effects could also be introduced by allowing citizens to move, for instance, in the two-dimensional plane in response to geographical and urban constraints, such as freeways, police stations or wealthier neighbourhoods, as well as in response to local strategy composition, giving rise to spatially dependent parameters $\epsilon(x, y)$, $\delta(x, y)$ and $\theta(x, y)$, for instance. We leave these possibilities as extensions of the current work.

In summary, due to the many elements included in this work – a heterogeneous society where players carry interdependent loss functions, varying recruitment costs and resource levels, memory effects in player payoffs – we find a rich array of potential optimal strategies, highlighting the fact that identifying optimal recruitment strategies is a non-trivial task. Perhaps the most representative case of a modern, western city in this work is that of a society with elevated resource levels and relatively low conversion costs. Here our findings are in agreement with general results from the sociological literature where it is argued that the major deterrent to crime is not necessarily more active enforcement or increasing the severity of punishment, but rather cooperation with authorities [2, 12]. Fighting crime is thus a collaborative effort of a community where engaging apathetics and villains could bring fruitful results.

Acknowledgements

This work was supported by the 'DyXi' project of the SYSCOMM programme of the French National Agency for Research through grant ANR-08-SYSC-008 (ABP); by the National Science Foundation through grants DMS-1021850 (MRD), DMS-0719642 (MRD) and DMS-0968309 (MBS); and through the ARO MURI grant W911NF-11-1-0332 (MBS and MRD).

Appendix A Necessary conditions for untargeted recruitment

In optimal control problems, the so-called 'Hamiltonian' is typically constructed, which is simply the integrand of the objective functional plus the dot product of a vector of adjoint variables and a vector obtained via the right-hand sides of the differential equations regulating the state variables. In the case of untargeted recruitment for our problem, this yields the following Hamiltonian $\mathscr{H}$:

$$\begin{aligned}\mathscr{H} = {} & (1-P-A)\big[(A+V)\delta + (1-A-V)(A+V)(\delta+\epsilon)\big] \\ & + \sum_{X\in\{P,A,V\}} \big[mC_X u_X X + \lambda_X \dot{X} + \lambda_{\ell_X}\dot{\ell}_X\big]. \end{aligned} \tag{A 1}$$

Here the $\dot{X}$ is defined as in equation (3.3) with the choice of untargeted recruitment, the $\dot{\ell}_X$ functions are given in equations (3.11)–(3.13) and the λ_X and λ_{ℓ_X} terms are the time-dependent adjoint variables.

According to Pontryagin's minimum principle [23], the optimal control variables $u_P^*(t)$, $u_A^*(t)$ and $u_V^*(t)$ are functions of time that minimize the above Hamiltonian in the

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

admissible region. However, because our Hamiltonian is linear in all of the control variables, the Hamiltonian is almost always minimized at the boundaries of the constraints in equation (3.2) – the so-called 'bang-bang' solution. That is, the optimal controls are given by

$$u_X^*(t) = \begin{cases} 1 & \text{if } \lambda_X(t) - C_X(t) > 0 \\ 0 & \text{if } \lambda_X(t) - C_X(t) < 0. \end{cases} \tag{A 2}$$

The solution is more ambiguous when $\lambda_X(t) - C_X(t) = 0$ for an extended period of time; these are referred to as singular optimal controls. To facilitate the numerical solution of the necessary conditions, especially in the presence of singular optimal controls, we follow [15] and add small quadratic control terms of the form

$$\frac{\xi}{2}[u_P(t)^2 + u_A(t)^2 + u_V(t)^2] \tag{A 3}$$

to the Hamiltonian. Using this new Hamiltonian, we find that the optimality condition

$$\frac{\partial \mathscr{H}}{\partial u_X} = mC_X X + \xi u_X - mX\lambda_X = 0 \tag{A 4}$$

is satisfied for

$$u_X^*(t) = \max\left\{0, \min\left[\frac{mX(t)\,[\lambda_X(t) - C_X(t)]}{\xi}, 1\right]\right\}. \tag{A 5}$$

In the case where $\xi \ll 1$, as is true in our numerical scheme, equation (A 5) gives essentially the same answer as equation (A 2) except in the case of a singular control.

The extension of Pontryagin's necessary conditions to systems with time delay is given in [9]. Using this result, the evolution equations of the adjoint variables for our problem are given by

$$\dot{\lambda}_X = \begin{cases} -\dfrac{\partial \mathscr{H}}{\partial X(t)} + \dfrac{\partial \mathscr{H}}{\partial X(t-\tau)}\Big|_{t+\tau} & \text{for } 0 \leqslant t \leqslant T - \tau \\ -\dfrac{\partial \mathscr{H}}{\partial X(t)}, & \text{for } t > T - \tau, \end{cases} \tag{A 6}$$

$$\dot{\lambda}_{\ell_X} = -\frac{\partial H}{\partial \ell_X(t)} \qquad \forall t, \tag{A 7}$$

which together with equation (3.3) (using untargeted recruitment), equations (3.11)–(3.13) and equation (A 5) form our Hamiltonian system with time delay.

The only item remaining to specify are boundary conditions. The initial conditions on the state variables are given by the history functions in equation (3.17) above. Since there are no final time conditions on the state variables, the adjoint variables must vanish at $t = T$:

$$\lambda_X(T) = \lambda_{\ell_X}(T) = 0. \tag{A 8}$$

These conditions are known as the natural boundary (or transversality) conditions that occur when the ends of the state variables are free [8, Section 2.3].

The numerical solution algorithm for our untargeted recruitment problem is based on the forward–backward sweep from [18], but adapted for delay differential equations. For

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

a given set of parameters, the solution is found as follows:

(1) Begin with $\xi \sim \mathcal{O}(1)$ and trial control functions $\hat{u}_X(t)$.

(2) Given the current values $\hat{u}_X(t)$, solve the system of equations (3.3) and (3.11)–(3.13) with initial conditions given in equation (3.17) for P, A, V, ℓ_P, ℓ_A and ℓ_V on the interval $[0, T]$.

(3) Using these solutions for the state variables, solve the adjoint equations (A 6) and (A 7) with terminal conditions given in equation (A 8) on the interval $[0, T]$.

(4) Calculate potential optimal controls $\hat{u}_X^*$ using equation (A 5).

(5) Check tolerance by comparing the current $\hat{u}_X$ to the potential optimal solutions $\hat{u}_X^*$. If the differences in these controls are smaller than a pre-determined tolerance, consider the solution (for the current value of ξ) found and proceed to the next step. If not, update $\hat{u}_X$ via $\hat{u}_X = q\hat{u}_X^* + (1-q)\hat{u}_X$, where $0 < q < 1$, and go back to step 2.

(6) If ξ is small enough that the terms of equation (A 3) are negligible (within tolerance) compared with the other terms of the objective functional, then $u_X^* = \hat{u}_X^*$ and the problem is solved. Otherwise, decrease ξ and go back to step 2.

It should be noted that the parameter q in step 5 may need to be quite small ($\approx 10^{-4}$) for the solution to converge. Hence, the code will dynamically adjust q downward if the solutions do not appear to be converging for the current value.

References

[1] Abrams, D. M. & Strogatz, S. H. (2003) Modelling the dynamics of language death. *Nature* **424**, 900.

[2] Akerlof, G. & Yellen, J. (1994) Gang behavior, law enforcement and community values. In: T. Taylor H. J. Aaron & T. E. Mann (editors), *Values and Public Policy*, Brookings Institution Press, NW Washington DC, pp. 173–209.

[3] Aletti, G., Naimzada, A. K. & Naldi, G. (2010) Mathematics and physics applications in sociodynamics simulation: The case of opinion formation. In: G. Naldi, L. Pareschi & G. Toscani (editors), *Mathematical Modeling of Collective Behavior in Socio-Economic and Life Sciences*, Birkhäuser, Boston, MA, pp. 203–221.

[4] Baker, T. E. (2005) *Introductory Crime Analysis, Crime Prevention and Intervention Strategies*, Pearson Custom Pub., Boston, MA.

[5] Baxter, G. J., Blythe, R. A., Croft, W. & McKane, A. J. (2006) Utterance selection model of language change. *Phys. Rev. E* **73**, 046118.

[6] Becerra, V. M. (2010) Solving complex optimal control problems at no cost with PSOPT. In: *2010 IEEE International Symposium on Computer-Aided Control System Design (CACSD) Part of the IEEE Multi-Conference on Systems & Control (MSC 2010)*, Piscataway, NJ, USA, IEEE, Washington, DC, pp. 1391–1396.

[7] Boccaletti, S., Latora, V., Moreno, Y., Chavez, M. & Huang, D. U. (2006) Complex networks: Structure and dynamics. *Phys. Rep.* **424**, 174–308.

[8] Bryson, A. E. & Ho, Y. C. (1969) *Applied Optimal Control: Optimization, Estimation, and Control*, Blaisdell, Waltham, MA.

[9] Budelis, J. J. & Bryson, A. E. (1970) Some optimal control results for differential-difference systems. *IEEE Trans. Autom. Control* **15**(2), 237–241.

[10] Castellano, C. (2005) Modeling cooperative behavior in the social sciences. In: J. Marro (editor), *AIP Conference Proceedings, Vol. 779*, American Institute of Physics, Melville, NY, 114 pp.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press

[11] Castellano, C., Fortunato, S. & Loreto, V. (2009) Statistical physics of social dynamics. *Rev. Mod. Phys.* **81**, 591–646.
[12] Ferrer, R. (2010) Breaking the law when others do: A model of law enforcement with neighborhood externalities. *Eur. Econ. Rev.* **54**, 163–180.
[13] Fossett, M. (2006) Ethnic preferences, social distance dynamics, and residential segregation: Theoretical explanations using simulation dynamics. *J. Math. Sociol.* **30**, 185–274.
[14] Helbing, D. & Lozano, S. (2010) Phase transitions to cooperation in the prisoner's dilemma. *Phys. Rev. Lett.* **81**, 057102.
[15] Jacobson, D. H., Gershwin, S. B. & Lele, M. M. (1970) Computation of optimal singular controls. *IEEE Trans. Autom. Control* **15**(1), 67–73.
[16] Kosmidis, K., Kalampokis, A. & Argyrakis, P. (2006) Statistical mechanical approach to human language. *Phys. A* **366**, 495–502.
[17] Le, S. & Boyd, R. (2007) Evolutionary dynamics of the continuous iterated prisoner's dilemma. *Phys. Rev. E* **245**, 258–267.
[18] Lenhart, S. & Workman, J. T. (2007) *Optimal Control Applied to Biological Models*, Chapman & Hall, London.
[19] Liggett, T. M. (2005) *Interacting Particle Systems*, Springer-Verlag, Berlin, Germany.
[20] Nardini, C., Kozma, B. & Barrat, A. (2008) Who's talking first? Consenus or lackthereof in coevolving opinion formation models. *Phys. Rev. Lett.* **100**, 158701.
[21] Newman, M. E. (2003) The structure and function of complex networks. *SIAM Rev.* **45**, 167–256.
[22] Newman, M. E., Watts, J. E. & Strogatz, S. H. (2002) Random graph models of social networks. *Proc. Natl. Acad. Sci. USA* **82**, 2566–2572.
[23] Pontryagin, L. S., Boltyanskii, V. G., Gamkrelize, R. V. & Mishchenko, E. F. (1962) *The Mathematical Theory of Optimal Processes*, Interscience, Hoboken, NJ.
[24] Rapoport, A. & Chammah, A. M. (1965) *Prisoner's Dilemma*. University of Michigan Press, Ann Arbor, MI.
[25] Santos, F. C. & Pacheco, J. M. (2005) Scale-free networks provide a unifying framework for the emergence of cooperation. *Phys. Rev. Lett.* **95**, 098104.
[26] Short, M. B., Brantingham, P. J. & D'Orsogna, M. R. (2010) Cooperation and punishment in an adversarial game: How defectors pave the way to a peaceful society. *Phys. Rev. E* **82**(6), 066114.1–066114.7.
[27] Sood, V. & Redner, S. (2005) Voter models on heterogeneous graphs. *Phys. Rev. Lett.* **94**, 178701.
[28] Stauffer, D. (2009) Opinion dynamics and sociophysics. In: R. Meyers (editor), *Encyclopedia of Complexity and Systems Science*, Springer, New York, pp. 6380–6388.
[29] Tyler, T. R. (2002) *Trust in the Law: Encouraging Public Cooperation with the Police and Courts*. Russel Sage Foundation, New York.
[30] Yizhaq, H., Portnov, B. A. & Meron, E. (2004) A mathematical model of segregation patterns in residential neighborhoods. *Env. Plan. A* **36**, 149–172.
[31] Zhang, Z. (2004) A dynamic model of residential segregation. *J. Math. Sociol.* **28**, 147–170.

https://doi.org/10.1017/S0956792512000332 Published online by Cambridge University Press