Graduated Punishments in Public Good Games

Allard van der Made *

IEEF, University of Groningen,
Nettelbosje 2, 9747AE Groningen, the Netherlands,
e-mail: `a.van.der.made@rug.nl`

November, 2012

Abstract

A host of social situations feature graduated punishments. We explain this phenomenon by studying a repeated public good game in which a social planner imperfectly monitors agents to detect shirkers. Agents' cost of contributing is private information and administering punishments is costly. A low punishment today imperfectly sorts agents by type: only low-cost agents contribute. The planner uses this information optimally by punishing tomorrow's (alleged) repeat shirkers harsher than first-time shirkers. The threat of becoming branded as repeat offender allows the planner to use a very mild punishment for first-time shirkers, attenuating the costs associated with administering punishments. Graduated punishments are consequently socially optimal as long as the population is not too homogeneous.

Keywords: graduated punishments, imperfect monitoring, collective action, reputation.

JEL classification codes: D82, H41, K49

*Financial support from the Belgian Federal Government through the IAP Project (contract 6/09) is gratefully acknowledged. This paper owes its existence to numerous discussions with Wouter Vergote. I would like to thank Remco van Eijkel, attendees at PET10, SMYE2011, EEA-ESEM 2011, and workshop participants at CEREC (Facultés universitaires Saint-Louis), at CORE (Université Catholique de Louvain), and at IEEF.

1 Introduction

A host of social situations involve collective action problems: from the point of view of the collective it is best if everybody acts in the interest of the group, yet it is individually optimal to act differently. Examples include tax avoidance, the tragedy of the commons, using polluting production technologies, and vote abstention. In many instances groups or societies have managed to induce individuals to behave in the interest of the collective. One important factor ensuring that individuals are inclined to choose the collectively preferred action is the presence of a monitoring institution that is able to punish (alleged) wrongdoers. Many scholars studying collective action problems have observed that successful punishment schemes often exhibit graduated sanctions: repeat offenders are punished more severely than first-time offenders (e.g. Agrawal, 2003, Ellickson, 1991, Ostrom, 1990, 2000, Wade, 1994). Graduated sanctions also appear in many judiciary systems, stipulating that habitual offenders can or must be punished more severely than first-time offenders.[1] In its most extreme form graduated punishments are such that first-time offenders receive a mere warning. Given its widespread use, it is surprising that this phenomenon has received limited theoretical attention.

We present a theory that explains the prevalence of punishment schemes featuring graduated punishments. We show that using graduated punishments is often optimal if monitoring is imperfect, administering punishments is costly, and agents differ with respect to how 'tempted' they are to choose the selfish action.

In our model a social planner faces a repeated public good problem. It is socially efficient if all agents contribute to the public good in each period, but an agent incurs a cost each time he contributes. The social planner monitors the behaviour of individual agents, but this monitoring is imperfect: some non-contributers (shirkers) escape being detected and some contributors are found guilty of shirking. The planner can administer punishments to alleged shirkers, but this is costly for society.[2] Moreover, because punishing an innocent person is in general seen as a grave injustice, we allow erroneously punishing a contributor to involve larger social costs than punishing a shirker.[3] The individually borne cost of contributing to the public good differs among agents and is either high or low. An agent's cost type is private information. The planner maximizes welfare, i.e. the social benefits of the contributions to the public good minus all costs.

Because punishing agents is costly, using a punishment that is sufficiently severe to deter all agents from shirking need not be optimal. Indeed, in a one-shot setting such a punishment is only optimal if the number of high-cost types is sufficiently large. If this number is not sufficiently large, then the social costs of erroneously administering severe

[1] For example, various state governments in the United States have enacted so-called *Three Strikes Laws*. Such laws require state courts to hand down a mandatory and extended period of incarceration to persons who have been convicted of a serious criminal offense on three or more separate occasions. See also `en.wikipedia.org/wiki/Three_strikes_law`.

[2] These costs include the administrative and legal costs associated with punishing someone. They can also include the cost of imprisoning someone for some time.

[3] See for instance the discussion in Chu et al. (2000, p. 130).

punishments to a large group of low-cost types outweigh the benefits of deterring a small group of high-cost types from shirking. The planner then sets a low punishment and only low-cost types contribute to the public good. If agents are not only supposed to contribute today, but also in future periods, then the planner can often improve upon the outcome of the one-shot setting by employing graduated punishments.

Using graduated punishments instead of a uniform punishment improves welfare for two reasons. Firstly, by imposing a mild sanction today the planner is able to (imperfectly) 'sort' agents by cost type: only low-cost types contribute if punishments are low. As a consequence, the planner can in the future (again imperfectly) tailor punishments to types by imposing a harsh punishment on repeat offenders and a moderate one on first-time offenders. This tailoring enables the planner to induce a given number of contributions in a more cost-efficient way.

Secondly, the mere threat of becoming 'branded' as shirker and thereby moving from the low-punishment regime to the high-punishment regime makes an agent reluctant to shirk today: since monitoring is imperfect and hence contributors are occasionally punished, being caught shirking today increases expected future punishments, even if the agent plans to contribute in future periods. In other words, an agent fears getting a reputation of being a shirker. This fear enables the planner to reduce the punishment for first-time shirkers below the low punishment of the one-shot setting (i.e. below the punishment that prevails if the number of high-cost types is small). This reputation effect is particularly strong if agents are patient. In fact, for all cost parameters one can find a discount rate above which it suffices to issue a mere warning to first-time offenders.

Using graduated punishments is not always optimal. If the society consists mainly of high-cost types, then using graduated punishments would yield a very low level of public good provision. To increase contributions the planner then opts for a uniform punishment that deters all agents from shirking, i.e. the high punishment of the one-shot setting. Nonetheless, because monitoring is imperfect, some agents are punished on the equilibrium path.[4] On the other hand, if the vast majority of agents incur the low cost when contributing, then most agents who end up in the high punishment regime are low-cost types. It can then be optimal to use the low punishment of the one-shot setting, as this leads to considerably lower punishment costs without significantly reducing the level of aggregate contributions.

Our results hinge crucially on the presence of type II errors, i.e. the possibility that the planner falsely judges someone guilty of shirking. If type II errors were completely absent, then only shirkers would be punished. The presence of type II errors has two effects. Firstly, if the planner would never erroneously punish contributors, then agents would not mind getting a bad reputation and the planner would consequently be unable to reduce the punishment administered to first-time shirkers below the low punishment of the one-shot setting. The reason that only type II errors matter in the determination of the reputation effect is that all agents contribute in each future period as soon as they move to the high-

[4] This is a common feature of equilibria of games with private information and imperfect monitoring. See e.g. Green and Porter (1984).

punishment regime. So, only type II errors lead to punishments being administered to agents who are branded as shirkers. Secondly, absent type II errors the planner always ensures that high-cost types contribute in the one-shot setting. The main advantage of setting a punishment that does not suffice to deter high-cost types from shirking is that such a low punishment entails low social costs of administering punishments to low-cost types. Yet, since a low-cost type is only punished if a type II error occurs, this advantage does not play a role if a contributor is never deemed guilty of shirking.[5]

If the planner knew each agent's cost type, then sorting agents by type would be redundant and the planner would therefore never resort to graduated punishments. With knowledge of agents' cost types she would be able to perfectly deter shirking: it suffices to 'promise' an agent an expected punishment at least as large as his cost of contributing.[6] So, the fact that an agent's cost of contributing is private information is essential for graduated punishments to arise.

Our framework not only applies to classic public good situations, but also to law enforcement problems. The cost of contributing to the public good is then replaced by the opportunity cost of not committing the crime under consideration. Furthermore, most crimes bestow a negative externality upon society at large. This ranges from commonly felt disgust following a gruesome murder to a reduction in the safety of online services caused by cybercrimes. Not engaging in criminal activities therefore increases welfare at the aggregate level in a similar fashion as contributing to a public good does.

Most collective action problems are plagued by limited monitoring possibilities and informational asymmetries. Consider for instance a groundwater basin shared by hundreds of farmers. Such basins can be destroyed by overextraction.[7] Whether a particular farmer extracts more water than he is entitled to is difficult to determine: a sudden drop in the water level could equally well be caused by overextraction by one of his neighbours. So, both type I and type II errors are bound to occur. How 'tempted' a farmer is to overextract water depends on unobservable psychological traits as well as the finer details of the microclimate and the soil composition he faces. His cost type is consequently private information.

This paper is organized as follows. Section 2 introduces the main ingredients of the model. The optimal punishment scheme of the one-shot setting can be found in Section 3. In Section 4 we study a two-periods version of our model. The infinite-horizon setting is analyzed in Section 5. In Section 6 we relate our work to the literature. Section 7 offers concluding remarks. All proofs are relegated to the Appendix.

[5]Type I errors reduce the probability that shirking is detected. Making a type I error with some probability ϵ_I has a similar impact on the optimal punishment(s) as only monitoring a sample containing a fraction $1-\epsilon_I$ of the population: a larger ϵ_I leads to higher actual punishments, but expected punishments remain constant.

[6]Since monitoring is imperfect, expected punishments are not equal to actual punishments.

[7]Ostrom (1990), chapter 4, gives an account of the collective action problems surrounding the groundwater basins near Los Angeles.

2 The Environment

A social planner faces a public good problem. If a fraction π of the population contributes to the public good, then the total social benefits of the public good amount to π. Contributing to the public good is costly. A fraction $1-\rho$ of the population consists of agents who incur the high cost γ when contributing, where $\gamma < 1$. The remaining fraction ρ consists of agents who incur the low cost $\gamma - \alpha$ when contributing, where $\alpha \in (0, \gamma)$. Because $\gamma < 1$, it is socially optimal if all agents contribute. Yet, for both high-cost and low-cost agents it is individually optimal to refrain from contributing, i.e. shirk. An agent's cost type (low-cost or high-cost) is private information. All agents are risk-neutral. We use the subscript L (H) to refer to low-cost (high-cost) types.

The planner can monitor agents' behaviour, enabling her to punish alleged shirkers. The planner's monitoring technology is flawed: with probability ϵ_I she fails to detect a shirker (a *type I error*) and with probability ϵ_{II} she erroneously judges someone guilty of shirking (a *type II error*). So, only a fraction $1-\epsilon_I$ of the shirkers are caught, whereas a fraction ϵ_{II} of the contributors are found guilty of something they did not do. We assume that monitoring agents is free, but that administering punishments is costly. Specifically, if the planner administers a punishment f, i.e. a punishment that reduces an agent's utility by f, then society bears a cost of cf, where $c > 0$. Furthermore, society bears an extra cost mf, where $m \geq 0$, when an innocent person is punished. So, the marginal social cost of punishing a shirker is c and the marginal social cost of punishing a contributor is $c+m$.

The planner maximizes welfare by choosing the punishments administered to alleged shirkers. These punishments are made public before agents advance to the contribution stage and we assume that the planner can commit to the announced punishments.[8] Welfare W consists of the social benefits of the public good, the individually borne costs of contributing, and the cost of administering punishments. In a one-shot setting welfare reads

$$W = \rho(1-\gamma+\alpha)\delta_L + (1-\rho)(1-\gamma)\delta_H - F, \tag{1}$$

where $\delta_L = 1$ ($\delta_L = 0$) if low-cost agents (do not) contribute, $\delta_H = 1$ ($\delta_H = 0$) if high-cost agents (do not) contribute, and F denotes the social costs of administering punishments. These costs amount to

$$F = \rho\big(\delta_L\epsilon_{II}(c+m) + (1-\delta_L)(1-\epsilon_I)c\big)f_0 + (1-\rho)\big(\delta_H\epsilon_{II}(c+m) + (1-\delta_H)(1-\epsilon_I)c\big). \tag{2}$$

where f_0 is the punishment that alleged shirkers face.

We assume that the laissez-faire outcome in which punishments are zero and no agent contributes is never optimal. To ensure that the planner never opts for laissez-faire we maintain the following condition throughout the paper:

Condition 1 *Laissez-faire is never optimal, specifically:* $1-\gamma > \frac{\epsilon_{II}}{1-\epsilon_I-\epsilon_{II}}(c+m)\gamma$.

[8] This assumption is not innocuous: because punishing agents is costly, ex post the planner prefers to refrain from punishing alleged shirkers. Assuming that the planner can commit to punishments is common practice in the literature: we share this assumption with, amongst others, Becker (1968).

If Condition 1 holds, then the planner prefers harsh punishments to laissez-faire, even if all agents are high-cost types ($\rho = 0$). The condition states that the gain in welfare $1-\gamma$ associated with a high-cost agent contributing must exceed the expected cost of incentivizing a high-cost agent to contribute by setting a sufficiently high punishment. This expected cost is the marginal cost $c+m$ of punishing a contributor times the probability ϵ_{II} of making a type II error times the required punishment $\frac{\gamma}{1-\epsilon_I-\epsilon_{II}}$.

Before we study settings with multiple periods, we derive the planner's optimal strategy in the one-shot setting. The one-shot outcome serves as a benchmark for the two-periods setting and the infinite-horizon setting we consider in Section 4 respectively Section 5: in these two settings the planner can always replicate the one-shot outcome in each period by simply using the optimal punishment of the one-shot setting. Our analysis of the one-shot setting therefore yields a lower bound on the per-period welfare that can be attained in the other settings.

3 The One-shot Setting

An agent contributes if the associated expected costs do not exceed the expected costs the agent faces when shirking.[9] A low-cost type consequently contributes if $\gamma - \alpha + \epsilon_{II} f_0 \leq (1-\epsilon_I) f_0$, i.e. if $f_0 \geq \frac{\gamma-\alpha}{1-\epsilon_I-\epsilon_{II}}$. A high-cost type contributes as long as $\gamma + \epsilon_{II} f_0 \leq (1-\epsilon_I) f_0$, which reduces to $f_0 \geq \frac{\gamma}{1-\epsilon_I-\epsilon_{II}}$. So, the planner chooses between the *low punishment* ($f_0 = \frac{\gamma-\alpha}{1-\epsilon_I-\epsilon_{II}}$) which only induces low-cost types to contribute and the *high punishment* ($f_0 = \frac{\gamma}{1-\epsilon_I-\epsilon_{II}}$) which ensures that high-cost types also contribute. Comparing the total welfare associated with the two possibilities yields

Proposition 1 *In the one-shot setting the social planner opts for*

$$\phi f_0^* = \begin{cases} \gamma & \text{if } \rho \leq \bar{\rho} \\ \gamma - \alpha & \text{if } \rho > \bar{\rho}, \end{cases} \tag{3}$$

where

$$\phi := 1 - \epsilon_I - \epsilon_{II} \tag{4}$$

measures the quality of the planner's monitoring technology and

$$\bar{\rho} := 1 - \frac{\frac{\epsilon_{II}}{\phi}(c+m)\alpha}{1-\gamma+c(\gamma-\alpha)-\frac{\epsilon_{II}}{\phi}m(\gamma-\alpha)} \in (0,1). \tag{5}$$

Since administering punishments is costly, it is not always optimal to induce all agents to contribute by using the high punishment. If the number of high-cost types is small (ρ large), then the increase in contributions brought about by moving from the low punishment to the high punishment is small. This move would also entail administering the high punishment instead of the low punishment to a fraction ϵ_{II} of the low-cost types. If the population

[9] We assume that an agent contributes if he is indifferent between contributing and shirking.

consists mainly of low-cost types the detrimental effect on welfare of administering a higher punishment to these agents dominates the positive effect of more contributions.

The social costs of erroneously punishing a fraction ϵ_{II} of the low-cost types are increasing in the probability ϵ_{II} that such a type II error occurs, the associated marginal social cost $c+m$, and the high individual cost γ. Ignoring the high-cost types by using the low punishment is particularly attractive if the associated social costs are small, which is the case if the low individual cost $\gamma - \alpha$ is small, i.e α is large, or if the marginal social cost of a type I error c is small. The planner therefore becomes more inclined to use the low punishment as ϵ_{II}, m, γ, or α increases. In other words, the threshold $\bar{\rho}$ is decreasing in these four parameters. Because c affects both the costs associated with erroneous punishments and those associated with just punishments, the impact of a change in c on $\bar{\rho}$ is ambiguous. As ϵ_I increases the difference between the two punishments grows, making the low punishment relatively more attractive. The threshold $\bar{\rho}$ therefore decreases in ϵ_I.

Observe that $\bar{\rho} \to 1$ as $\epsilon_{II} \to 0$. So, the planner always opts for the high punishment $\frac{\gamma}{\phi}$ if type II errors are never made. The reason is that the main advantage of using the low punishment $\frac{\gamma-\alpha}{\phi}$ disappears as ϵ_{II} approaches 0: A lower punishment entails lower social costs of erroneously administering punishments to low-cost types. Yet, since these agents are only punished if a type II error occurs, this advantage is absent if $\epsilon_{II} = 0$.

The trade-off between higher contributions and lower social costs also plays a role in a setting with two periods. Yet, it turns out that the planner often uses information regarding agents' past behaviour to alleviate the social costs of punishments.

4 The Two-periods Setting

Agents are now supposed to contribute to the public good twice: in period 1 and in period 2. An agent's type is again private information. The planner recalls in period 2 whether or not she has punished a given agent in period 1. Just like in the one-shot setting the planner makes a type i error with probability ϵ_i when investigating an agent's behaviour, $i \in \{I, II\}$. Drawing the wrong conclusion regarding a particular agent's behaviour in period 1 does not affect the probability with which she misjudges that agent's behaviour in period 2.

Recalling who has been punished in period 1 enables the planner to use differentiated punishments in period 2, one for agents who have not been punished in period 1 (f_2) and one for agents who have been punished ($\hat{f}_2$). The planner can only use one punishment (f_1) in period 1. The planner announces all punishments at the start of the game.[10] Each agent employs backward induction to arrive at his optimal strategy. The timing of the game is as follows:

0. The planner announces the punishments.

1a. Each agent decides whether to contribute or to shirk (δ_L and δ_H are chosen).

[10] We again assume that the planner can commit to the announced punishments.

1b. The planner carries out investigations and administers punishments.

2a. Each agent decides whether to contribute or to shirk.

2b. The planner carries out investigations and administers punishments.

Payoffs are realized at the end of each period. An agent minimizes his total expected costs. The planner maximizes total welfare $\mathcal{W}$, the sum of welfare in period 1 (W_1) and welfare in period 2 (W_2).

Whether the planner does use differentiated punishments is the subject of the next subsection.

4.1 Analysis

Using two different punishments in period 1 can only be optimal if low-cost types and high-cost types behave differently in period 1. The reason is that the planner cannot distill any information regarding an agent's type from his behaviour in period 1 if the two types employ the same strategy in that period. So, the planner only uses differentiated punishments if $\delta_L = 1$ and $\delta_H = 0$.[11]

Since the game ends after period 2, an agent contributes in period 2 if the associated expected cost does not exceed the expected cost associated with shirking. If the planner contemplates using two different punishments, she can therefore confine attention to $f_2^* = \frac{\gamma-\alpha}{\phi}$ (for those who have not been punished in period 1) and $\hat{f}_2^* = \frac{\gamma}{\phi}$ (for those who have been punished in period 1). The pair $(f_2^*, \hat{f}_2^*)$ induces high-cost types who are punished in period 1 as well as all low-cost types to contribute in period 2. Because the expected punishment ϕf_2^* is less than their cost of contributing γ, high-cost types who dodged being punished in period 1 shirk again in period 2 when faced with this pair of punishments.[12]

The planner induces the period 1-choices $\delta_L = 1$ and $\delta_H = 0$ by setting a moderate punishment f_1 that abides by the following incentive compatibility constraints:

- Low-cost types prefer to contribute in period 1 if:

$$\gamma - \alpha + \epsilon_{II}(f_1 + \gamma - \alpha + \epsilon_{II}\hat{f}_2^*) + (1-\epsilon_{II})(\gamma - \alpha + \epsilon_{II} f_2^*) \leq (1-\epsilon_I)(f_1 + \gamma - \alpha + \epsilon_{II}\hat{f}_2^*) + \epsilon_I(\gamma - \alpha + \epsilon_{II} f_2^*).$$

  The left-hand side of this constraint consists of the expected costs a low-cost type incurs when contributing in period 1. It equals the cost of contributing $\gamma - \alpha$ plus the expected costs associated with being erroneously punished in period 1 and/or period 2. The right-hand side consists of the expected costs a low-cost type faces

[11] Because low-cost types incur a lower cost than high-cost types when contributing, we can immediately discard the possibility that $\delta_L = 0$ and $\delta_H = 1$.

[12] Expressions like ϕf_2^* actually denote a *difference* in expected punishments: $\phi f_2^* = (1-\epsilon_I) f_2^* - \epsilon_{II} f_2^*$ is the expected punishment faced when shirking minus the expected punishment faced when contributing. We omit the "difference in" for ease of exposition.

when shirking in period 1. Note that we have used the fact that low-cost types always contribute in period 2 if the pair $(f_2^*, \hat{f}_2^*)$ is used. Using the expressions for f_2^* and $\hat{f}_2^*$ and (4) reduces the constraint to

$$\phi f_1 \geq \gamma - \alpha - \epsilon_{II}\alpha. \tag{6}$$

- High-cost types prefer to shirk in period 1 if:

  $$\gamma + \epsilon_{II}(f_1 + \gamma + \epsilon_{II}\hat{f}_2^*) + (1-\epsilon_{II})(1-\epsilon_I)f_2^* > (1-\epsilon_I)(f_1 + \gamma + \epsilon_{II}\hat{f}_2^*) + \epsilon_I(1-\epsilon_I)f_2^*.$$

  The left-hand side of this constraint consists of the expected costs a high-cost type incurs when contributing in period 1 and the right-hand side consists of that type's expected costs when shirking in period 1. Observe that we have used the fact that a high-cost type only contributes in period 2 if he is punished in period 1. Rewriting the constraint yields

  $$\phi f_1 < \gamma - \alpha + \epsilon_I\alpha. \tag{7}$$

The incentive compatibility constraints (6)-(7) reveal that if the planner opts for differentiated punishments in period 2, then she sets

$$\phi f_1^* = \max\{\gamma - \alpha - \epsilon_{II}\alpha, 0\}.$$

It remains to determine when differentiated punishments are optimal. The *menu of punishments* $\boldsymbol{f}^* := (f_1^*, f_2^*, \hat{f}_2^*)$ is optimal if the total welfare $\mathcal{W}(\boldsymbol{f}^*)$ it generates exceeds the total welfare $\mathcal{W}(f_0^*)$ society enjoys should the planner use the single punishment f_0^* given in (3) in both periods. Comparing these two welfare expressions results in

Proposition 2 *There exist $\check{\rho} \in (0, \bar{\rho})$ and $\hat{\rho} > \bar{\rho}$ such that the social planner maximizes total welfare by using the menu of punishments*

$$f_1^* = \max\{\tfrac{\gamma-\alpha}{\phi} - \tfrac{\epsilon_{II}\alpha}{\phi}, 0\}, \quad f_2^* = \tfrac{\gamma-\alpha}{\phi}, \quad \hat{f}_2^* = \tfrac{\gamma}{\phi} \tag{8}$$

if $\rho \in (\check{\rho}, \hat{\rho})$. If either $\rho < \check{\rho}$ or $\rho > \hat{\rho}$, then the social planner maximizes total welfare by using the single punishment f_0^* given in (3) in both periods. The upper bound $\hat{\rho}$ equals 1 if and only if $\gamma - \alpha - \epsilon_{II}\alpha \geq 0$.

If $\boldsymbol{f}^*$ is used, then low-cost types always contribute whereas a high-cost type shirks in period 1 and contributes in period 2 only if he has been punished in period 1. If the population is (all but) homogeneous (ρ close to 0 or 1), then the planner opts for the single punishment given in (3) instead of graduated punishments. This is intuitive: when faced with a homogeneous population the best the planner can do is to use the smallest punishment that deters agents of the extant type from shirking in both periods.

The positive effects of using graduated punishments (i.e. of using $\boldsymbol{f}^*$) instead of a single punishment start playing a role as ρ departs from 0 or 1. If the population is heterogeneous, then using graduated punishments allows the planner to imperfectly sort agents by type.

The reason is that the period 1-punishment f_1^* is too low to incentivize high-cost types to contribute and hence only low-cost types contribute in period 1. This implies that an alleged period 2-shirker who has already been punished in period 1 is likely to be a high-cost type. Because the planner occasionally draws the wrong conclusion when investigating agents' behaviour, this mechanism only imperfectly sorts agents by type.

This sorting enables the planner to tailor period 2-punishments to types to a large extent. Since the vast majority of those who have been found guilty of shirking in period 1 are high-cost types and such types can only be deterred from shirking by 'promising' them an expected punishment of at least γ, the planner uses the punishment $\frac{\gamma}{\phi}$ for repeat offenders. On the other hand, an agent who is found guilty of shirking for the first time in period 2 is probably a low-cost type. The punishment $\frac{\gamma-\alpha}{\phi}$ therefore suffices to deter most of the agents who were not punished in period 1 from shirking in period 2.

Both the low punishment $\frac{\gamma-\alpha}{\phi}$ and the high punishment $\frac{\gamma}{\phi}$ of the one-shot setting come with a severe drawback: the low punishment leads to a suboptimal contribution level (only a fraction ρ of the population contributes) whereas the high punishment leads to considerable social costs of administering punishments. If the planner can tailor punishments to types, albeit imperfectly, then the planner does not face a choice between two severe drawbacks. Firstly, with graduated punishments only high-cost types who escaped being punished in period 1 shirk in period 2 and the contribution level in period 2 consequently exceeds ρ. Secondly, by only administering the high punishment $\frac{\gamma}{\phi}$ to repeat offenders, the planner moderates the social costs of administering punishments.

If the population consists mainly of high-cost types (ρ small), then using graduated punishments would result in a very low level of public good provision in period 1. At the same time a considerable part of the population, namely a fraction $1-\epsilon_I$ of the high-cost types, would be punished in that period. The associated social costs become smaller as ρ increases: the number of period 1-shirkers and hence the number of agents who receive the period 1-punishment decreases in ρ. Furthermore, the level of public good provision in period 1 increases in ρ. So, the downsides of using graduated punishments become less severe as ρ increases. This explains why the lower bound $\check{\rho}$ always exceeds 0 whereas the upper bound $\hat{\rho}$ often equals 1.

Observe that the period 1-punishment $f_1^* = \frac{\gamma-\alpha-\epsilon_{II}}{\phi}$ is less than $\frac{\gamma-\alpha}{\phi}$, the smallest punishment that deters low-cost types from shirking in the one-shot setting. The reason that the planner is able to incentivize low-cost types to contribute in period 1 with an expected punishment below their cost of contributing $\gamma-\alpha$ is that an agent found guilty of shirking in period 1 receives part of his 'effective punishment' indirectly. Such an agent not only faces the (direct) punishment f_1^*, but he will also receive the high punishment $\hat{f}_2^*$ instead of the lower punishment f_2^* should he be found guilty of shirking a second time. So, the threat of becoming known as a repeat offender, i.e. the fear of getting a bad reputation, allows the planner to reduce the expected punishment used in period 1 below the low punishment that is required in a one-shot setting. The size of this *reputation effect* equals the loss in expected utility stemming from getting a bad reputation: $\frac{\epsilon_{II}\alpha}{\phi}$ is the difference between $\hat{f}_2^*$ and f_2^* times the probability that a contributing agent is erroneously found guilty of shirking. If $\gamma - \alpha - \epsilon_{II}\alpha \leq 0$, then $f_1^* = 0$. So, if the probability that the planner makes a type II error is sufficiently large, then she merely *warns* an alleged period 1-shirker. In that case the size of the reputation effect is smaller than $\frac{\epsilon_{II}\alpha}{\phi}$.[13]

5 The Infinite-horizon Setting

Time $t = 1, 2, 3, \ldots$ is discrete. Each period t consists of three stages. In the first stage each agent chooses between contributing to the public good and shirking. In the second stage the social planner carries out investigations and punishes agents who have been found guilty of shirking. In the last stage, the *renewal stage*, a fraction $1 - \beta \in (0, 1)$ of the population dies and is replaced by new agents. The probability that a given agent dies does not depend on his type, how often he has shirked, or the number of times he has been punished. So, each agent advances to the next period with probability β. The population is again characterized by the parameters γ, α, and ρ. In particular, a fraction ρ of each generation and thus of the population in any period incurs the low cost $\gamma - \alpha$ when contributing. Since $\gamma < 1$, it is socially optimal that an agent contributes in each period that he lives.

The planner is immortal and keeps track of whether or not a given agent has been punished in the past. The quality of her monitoring technology is again characterized by the probabilities ϵ_I and ϵ_{II}. She announces all punishments that might be applicable in period t at the start of that period, before agents decide whether to contribute or to shirk. She can opt to use two different punishments, one for agents who have never been punished before (f_t) and one for agents who have been punished at least once ($\hat{f}_t$). Alternatively, she can administer the same punishment to all alleged shirkers. Since the fraction low-cost types is ρ in each period, the planner chooses the punishment f_0^* given in (3) in the latter case.

The population can be divided in four categories: low-cost types who have never been punished, low-cost types who have been punished at least once, high-cost types who have never been punished, and high-cost types who have been punished at least once. We focus on the *stationary equilibria* of the model, i.e. equilibria that can prevail if the composition of the population with respect to the above categorization remains unaltered as the economy moves from some period to the next one. So, we focus on the very long run ($t \to \infty$).

In each period the planner aims to maximize the welfare generated in that period. Observe that a strategy of the planner that supports a stationary equilibrium in the present setting also supports the corresponding stationary equilibrium of the game in which the planner maximizes current welfare *plus* discounted future welfare, irrespective of the dis-

[13] If $f_1^* = \frac{\gamma - \alpha - \epsilon_{II}\alpha}{\phi}$ and $\rho = 1$, then the reduction in aggregate punishments in period 1 due to the reputation effect equals the increase in aggregate punishments in period 2 caused by the fact that some agents receive the punishment $\hat{f}_2^*$ instead of f_2^*. The planner therefore becomes indifferent between using f_0^* and using $\boldsymbol{f}^*$ as $\rho \to 1$ if $\gamma - \alpha - \epsilon_{II}\alpha \geq 0$, i.e. if the size of the reputation effect is $\frac{\epsilon_{II}\alpha}{\phi}$. If the reputation effect is smaller, then it does not fully offset the detrimental effect on welfare of administering the punishment $\hat{f}_2^*$ instead of f_2^* to repeat offenders and $\hat{\rho}$ is consequently smaller than 1.

count rate. An agent minimizes his expected current and discounted future costs. The only difference between the planner and agents regarding their attitude towards the future stems from the fact that the former is immortal whereas an agent dies with probability $1-\beta$ at the end of a period. It is therefore natural to use β as the agents' discount factor between periods. Payoffs are realized after the punishment stage, but before the renewal stage.

A stationary equilibrium is supported by a pair of punishments f^* and $\hat{f}^*$ and four contribution rules: δ_L^*, $\hat{\delta}_L^*$, δ_H^*, and $\hat{\delta}_H^*$. Here, $\delta_j^* = 1$ ($\delta_j^* = 0$) if a type j-agent who has never been punished decides (not) to contribute, $j = L, H$. Similarly, $\hat{\delta}_j^* = 1$ ($\hat{\delta}_j^* = 0$) if a type j-agent who has been punished at least once decides (not) to contribute, $j = L, H$. Of course, the equilibrium contribution rules must be best responses to the equilibrium punishments and vice versa.

The optimal strategy of the planner depends on the composition of the population. In the next subsection we first derive the composition of the population as $t \to \infty$ before we determine the stationary equilibria of the game. We omit any reference to taking limits in these subsections if there is no risk of confusion.

5.1 Analysis

Whether a given agent has been punished in the past is irrelevant if the planner opts for the uniform punishment given in (3). Just like in the two-periods setting using graduated punishments can only be optimal if these punishments are such that low-cost types always contribute whereas a high-cost type only contributes if he has been punished at least once. We can thus confine attention to the contribution strategies $(\delta_L, \hat{\delta}_L, \delta_H, \hat{\delta}_H) = (1, 1, 0, 1)$.

Let $\hat{q}$ (q) be the fraction of the population that has (never) been punished in the past. Denote the fraction of the population that consists of low-cost types who are in q ($\hat{q}$) by μ ($\hat{\mu}$).[14] By definition $\hat{q} = 1 - q$ and $\hat{\mu} = \rho - \mu$. Furthermore, the fraction of the population that consists of high-cost types who are in q equals $q - \mu$. Using the facts that a fraction $1 - \delta$ of the old population is replaced by new agents and that $(\delta_L, \delta_H) = (1, 0)$ one infers that in a stationary equilibrium with graduated punishments q abides by the following 'flow equation':

$$\begin{aligned} q =& (1-\beta) + \beta\mu\big(\delta_L(1-\epsilon_{II}) + (1-\delta_L)\epsilon_I\big) + \beta(q-\mu)\big(\delta_H(1-\epsilon_{II}) + (1-\delta_H)\epsilon_I\big) \\ =& 1 - \beta + \beta\phi\mu + \beta\epsilon_I q. \end{aligned}$$

The right-hand side of this equation contains the inflow of new agents (which equals $1-\beta$) and the agents who stay in q because they have not been punished in the previous period (which equals a fraction $1 - \epsilon_{II}$ of the contributors in q plus a fraction ϵ_I of the shirkers in q). The flow equation for μ reads

$$\mu = (1-\beta)\rho + \beta\mu\big(\delta_L(1-\epsilon_{II}) + (1-\delta_L)\epsilon_I\big) = (1-\beta)\rho + \beta(1-\epsilon_{II})\mu.$$

[14]We allow ourselves a slight abuse of notation by using $\hat{q}$ (q) for the fraction of the population that has (never) been punished in the past as well as for the *set* of agents with this feature.

Combining the two flow equations yields

Lemma 1 *Suppose that $(\delta_L, \delta_H) = (1, 0)$. Then:*

$$q = \frac{(1-\beta)(1-\beta\epsilon_I - \beta\phi(1-\rho))}{(1-\beta\epsilon_I)(1-\eta)}, \quad \mu = \frac{(1-\beta)\rho}{1-\eta}, \tag{9}$$

where $\eta := \beta(1-\epsilon_{II})$ is the probability that a contributor stays in q for one more period.

Observe that

$$\frac{\mu}{q} = \frac{1-\beta\epsilon_I}{1-\beta\epsilon_I - \beta\phi(1-\rho)} \times \rho > \rho,$$

i.e. in q the fraction of low-cost types exceeds ρ. This is intuitive: because low-cost types in q contribute, most of them stay in q. On the other hand, the majority of the high-cost types, being found guilty of shirking, move to $\hat{q}$ and hence $\frac{\hat{q}-\hat{\mu}}{\hat{q}} = 1 - \frac{\hat{\mu}}{\hat{q}} > 1-\rho$.

Let us now determine for which pair of punishments f and $\hat{f}$ the contribution strategies $(\delta_L, \hat{\delta}_L, \delta_H, \hat{\delta}_H) = (1, 1, 0, 1)$ prevail. An agent minimizes his expected discounted costs by choosing between contributing and shirking.[15] Denote the *continuation cost* of a type j-agent who is in q ($\hat{q}$) by C_j ($\hat{C}_j$), $j = L, H$. Then agents' behaviour is governed by the following four Bellman equations:

- Bellman equation for low-cost types who have never been punished:

$$C_L = \min_{\delta_L \in \{0,1\}} \Big[\delta_L\big(\gamma - \alpha + \epsilon_{II}(f + \beta\hat{C}_L) + (1-\epsilon_{II})\beta C_L\big) + (1-\delta_L)\big((1-\epsilon_I)(f+\beta\hat{C}_L) + \epsilon_I\beta C_L\big)\Big]. \tag{10}$$

  If a low-cost type in q contributes ($\delta_L = 1$), then he incurs the cost $\gamma - \alpha$. With probability ϵ_{II} he is erroneously found guilty of shirking in which case he receives the punishment f and moves to $\hat{q}$. If he is not punished, which happens with probability $1 - \epsilon_{II}$, then he stays in q. To understand the $(1-\delta_L)$-part of (10), note that the planner detects shirking with probability $1-\epsilon_I$, in which case the agent receives the punishment f and moves to $\hat{q}$. With probability ϵ_I the shirking agent escapes being punished and stays in q. In all cases the agent advances to the next period with probability β.

- Bellman equation for low-cost types who have been punished in the past:

$$\hat{C}_L = \min_{\hat{\delta}_L \in \{0,1\}} \Big[\hat{\delta}_L\big(\gamma - \alpha + \epsilon_{II}\hat{f}\big) + (1-\hat{\delta}_L)(1-\epsilon_I)\hat{f} + \beta\hat{C}_L\Big]. \tag{11}$$

[15] If low-cost types or high-cost types would employ a mixed strategy, then an infinitesimal increase in (one of) the punishment(s) would lead to a discrete upward jump in contributions. This renders mixed strategy equilibria impossible. We can thus confine attention to pure strategies.

Since an agent cannot escape from $\hat{q}$, C_L is absent from (11). Furthermore, if an agent in $\hat{q}$ is found guilty of shirking, which happens with probability ϵ_{II} if that agent contributes and with probability $1-\epsilon_I$ if that agent shirks, then he receives the punishment $\hat{f}$.

- Bellman equation for high-cost types who have never been punished:

$$C_H = \min_{\delta_H \in \{0,1\}} \Big[\delta_H\big(\gamma+\epsilon_{II}(f+\beta\hat{C}_H)+(1-\epsilon_{II})\beta C_H\big)+(1-\delta_H)\big((1-\epsilon_I)(f+\beta\hat{C}_H)+\epsilon_I\beta C_H\big)\Big]. \tag{12}$$

  This equation is akin to (10): the main difference stems from the fact that a high-cost type incurs the cost γ instead of the cost $\gamma-\alpha$ when he contributes.

- Bellman equation for high-cost agents who have been punished in the past:

$$\hat{C}_H = \min_{\hat{\delta}_H \in \{0,1\}} \Big[\hat{\delta}_H\big(\gamma+\epsilon_{II}\hat{f}\big)+(1-\hat{\delta}_H)(1-\epsilon_I)\hat{f}+\beta\hat{C}_H\Big]. \tag{13}$$

  Construction of this equation mirrors that of (11).

One easily verifies that $\hat{\delta}_L = 1$ is optimal if $\phi\hat{f} \geq \gamma-\alpha$ and that $\hat{\delta}_H = 1$ is optimal if $\phi\hat{f} \geq \gamma$. Consequently, if the planner does use differentiated punishments, then $\hat{f} = \frac{\gamma}{\phi}$. With this punishment for repeat offenders and given the contribution strategies $\hat{\delta}_L = \hat{\delta}_H = 1$ the continuation costs for agents in $\hat{q}$ become

$$\hat{C}_L\big|_{\hat{f}=\frac{\gamma}{\phi}} = \frac{\gamma-\alpha+\epsilon_{II}\frac{\gamma}{\phi}}{1-\beta}, \quad \hat{C}_H\big|_{\hat{f}=\frac{\gamma}{\phi}} = \frac{\gamma+\epsilon_{II}\frac{\gamma}{\phi}}{1-\beta}. \tag{14}$$

These continuation costs equal the discounted costs of contributing in each period plus the discounted expected (erroneous) punishments.

From (10) one gathers that low-cost types in q contribute if $\gamma-\alpha \leq \phi f+\phi\beta(\hat{C}_L-C_L)$. From (12) it follows that high-cost types in q shirk if $\gamma > \phi f+\phi\beta(\hat{C}_H-C_H)$. These two inequalities are the incentive compatibility constraints that must hold if the planner opts for graduated punishments. Combining these constraints with (14) yields[16]

$$\phi f \geq \gamma-\alpha-\tfrac{\beta}{1-\beta}\epsilon_{II}\alpha, \quad \phi f < \gamma. \tag{15}$$

The incentive compatibility constraint for low-cost types ($\phi f \geq \gamma-\alpha-\frac{\beta}{1-\beta}\epsilon_{II}\alpha$) resembles its counterpart in the two-periods setting (see (6)). The only difference between the two constraints is that the reduction in the required punishment stemming from the reputation effect is now multiplied by $\frac{\beta}{1-\beta}$. This number is an agent's life expectancy and hence an agent who plans to contribute in each period expects to be erroneously

[16] See the Appendix for details.

punished $\frac{\beta}{1-\beta}\epsilon_{II}$ times.[17] The incentive compatibility constraint for high-cost types does differ dramatically from its counterpart (7). The reason is that the planner uses two punishments if she opts for graduated punishments in the infinite-horizon setting, whereas she uses three punishments (the menu $\boldsymbol{f}^*$) is he opts for graduated punishments in the two-periods setting. Since both the expected punishment for agents in $\hat{q}$ and a high-cost type's cost of contributing are equal to γ, a high-cost type shirks when in q as long as the expected punishment for agents in q is less than γ.

Since $\gamma-\alpha-\frac{\beta}{1-\beta}\epsilon_{II}\alpha < \gamma$, the planner can always find a punishment f such that agents opt for $(\delta_L, \delta_H) = (1, 0)$ given that $\hat{f} = \frac{\gamma}{\phi}$. Inducing these contribution strategies is often optimal:

Proposition 3 *There exist $\check{\rho}_\infty \in (0, \bar{\rho})$ and $\hat{\rho}_\infty \in (\bar{\rho}, 1]$ such that the social planner maximizes per-period welfare by using the pair of punishments*

$$f^* = \max\{\tfrac{\gamma-\alpha}{\phi} - \tfrac{\beta}{1-\beta}\tfrac{\epsilon_{II}\alpha}{\phi}, 0\}, \quad \hat{f}^* = \tfrac{\gamma}{\phi} \tag{16}$$

if $\rho \in (\check{\rho}_\infty, \hat{\rho}_\infty)$. If either $\rho < \check{\rho}_\infty$ or $\rho > \hat{\rho}_\infty$, then the social planner maximizes per-period welfare by using the single punishment f_0^* given in (3) in both periods. The upper bound $\hat{\rho}_\infty$ equals 1 if and only if $\gamma - \alpha - \frac{\beta}{1-\beta}\epsilon_{II}\alpha \geq 0$.

If graduated punishments are used, then low-cost types always contribute whereas high-cost types shirk as long as they have not yet received a punishment. A high-cost type shirks on average $\frac{1}{1-\beta\epsilon_I}$ times.[18]

Recall that in the two-periods setting agents only fear getting a bad reputation in the first period. The incentives of agents consequently differ across periods and the planner therefore has to use three different punishments when opting for graduated punishments: one for those found guilty of shirking in period 1, one for first-time offenders in period 2, and one for repeat offenders. By contrast, only two punishments are used in the stationary equilibrium of the infinite-horizon setting. The reason is that getting a bad reputation always increases an agent's expected discounted future costs in the infinite-horizon setting. In that setting the planner can therefore always administer the cost-efficient punishment f^* to alleged first-time offenders.

The (maximal) difference between the punishment for first-time offenders f^* and the low punishment of the one-shot setting ($\frac{\gamma-\alpha}{\phi}$), i.e. $\frac{\beta}{1-\beta}\frac{\epsilon_{II}\alpha}{\phi}$, equals the size of the reputation effect of the two-periods setting ($\frac{\epsilon_{II}\alpha}{\phi}$) times an agent's life expectancy ($\frac{\beta}{1-\beta}$). So, the size of the reputation effect in the infinite-horizon setting is the loss in expected per-period utility stemming from getting a bad reputation times the expected number of periods that an agent stays alive. As β becomes sufficiently large the planner arrives at a corner solution

[17] An agent stays alive for exactly k periods after the current period with probability $\beta^k(1-\beta)$. His life expectancy thus equals $\sum_{k\in\mathbb{N}} k\beta^k(1-\beta) = \beta(1-\beta)\frac{\mathrm{d}}{\mathrm{d}\beta}\left(\sum_{k\in\mathbb{N}}\beta^k\right) = \frac{\beta}{1-\beta}$.

[18] With probability $1-\epsilon_I+\epsilon_I(1-\beta) = 1-\beta\epsilon_I$ a shirking high-cost type is caught shirking or fails to advance to the next period. In both cases he stops shirking. With the complementary probability $\beta\epsilon_I$ he advances to the next period and shirks in that period. So, the expected number of times a high-cost type shirks is $\sum_{k=0}^{\infty}(k+1)(1-\beta\epsilon_I)(\beta\epsilon_I)^k = \frac{1}{1-\beta\epsilon_I}$.

in which she merely issues warnings to first-time offenders ($f^* = 0$). The reason that warnings suffice to induce low-cost types to contribute is intuitive: the larger β is, the more important expected future costs are (relative to costs incurred in the current period) and the more agents fear moving to $\hat{q}$ and the lower f^* consequently can be.

Even though a solution with warnings entails zero costs of punishing first-time offenders, it is suboptimal if the fraction of low-cost types ρ is very large. The rationale behind this result has already been alluded to in footnote ..: If $\rho = 1$ and $f^* = \frac{\gamma-\alpha}{\phi} - \frac{\beta}{1-\beta}\frac{\epsilon_{II}\alpha}{\phi}$, then the reduction in aggregate punishments to alleged first-time offenders due to the reputation effect equals the increase in aggregate punishments to alleged repeat offenders due to the fact that they receive the punishment $\frac{\gamma}{\phi}$ instead of $\frac{\gamma-\alpha}{\phi}$. However, if $\frac{\gamma-\alpha}{\phi} < \frac{\beta}{1-\beta}\frac{\epsilon_{II}\alpha}{\phi}$, then the planner, being forced to set $f^* = 0$, cannot fully exploit the reputation effect and the reduction in aggregate punishments to alleged first-time offenders is consequently smaller than the increase in aggregate punishments to alleged repeat offenders.

Note that if both ρ and β are close to 1, then the population consists mainly of (long-lived) low-cost types. Because β is large, it is very likely that such a low-cost type spends a large part of his life in $\hat{q}$: since monitoring is imperfect, the probability that a law-abiding agent is found guilty of shirking at least once in τ periods goes to 1 as $\tau \to \infty$. In fact, $\hat{q}$, the fraction of the population that has been punished at least once, converges to 1 as $\beta \uparrow 1$. The vast majority of the agents in $\hat{q}$ are thus low-cost types if ρ and β are both large. Administering the high punishment $\hat{f}^*$ to low-cost types is clearly suboptimal: the punishment $\frac{\gamma-\alpha}{\phi}$ suffices to deter these agents from shirking. Because the number of high-cost types in $\hat{q}$ is negligible if ρ is close to 1, administering the punishment f^* to repeat offenders is dominated by administering the more cost-efficient punishment $\frac{\gamma-\alpha}{\phi}$. The planner therefore does not use graduated punishments if ρ and β are both close to 1.

6 Relation to the Literature

Graduated punishments have received quite some theoretical attention, most notably from law and economics scholars. Various explanations for this phenomenon have been proposed. Miceli and Bucci (2005) argue that the dire labour market prospects of convicted criminals makes committing crimes relatively more attractive for those who already have a criminal record. This effect can be negated by punishing repeat offenders harsher than first-time offenders. If offenders learn how to evade apprehension, as in Mungan (2010), then the *expected* punishment a repeat offender faces is lower than the expected punishment a first-time offender faces should the *actual* punishment remain the same. It is then optimal to set the actual punishment for repeat offenders higher than the actual punishment for first-time offenders. Of course, law enforcers could also learn from past offenses, yielding an increase in the probability that repeat offenses are detected. If law enforcers learn more than offenders, then the optimal punishment for repeat offenders is lower than the one for first-time offenders.[19]

[19]Dana (2001) provides ample arguments in favour of a higher probability of detection for repeat offenders.

Stigler (1970) argued informally that heavy penalties are unnecessary for first-time offenders if they are likely to have committed the offense accidentally and the probability of repetition is negligible. In Rubinstein (1979) offenses may also have been committed by accident. Convicting innocent offenders is detrimental to welfare. Rubinstein shows that it is then optimal to be lenient towards individuals with a 'reasonable' criminal record, i.e. those individuals are not administered the exogenously given punishment. Erroneous convictions also play a central role in Chu et al. (2000). Their planner tries to minimize total social costs, which consists of the harm imposed on society by criminal conduct and the cost of erroneous convictions. Chu et al. (2000) establish that in a two-period setting society is always best off if alleged repeat offenders are punished more severely than alleged first-time offenders. Such a solution is optimal, because the probability of convicting an innocent offender twice is much lower than convicting an innocent offender only once. Since punishing those who did commit crimes is costless, Chu et al.'s planner does not face a trade-off between crime prevention and cost minimization comparable to our trade-off between public good provision and cost minimization. Furthermore, they do not allow the punishment for first-time offenders in period 1 to differ from its counterpart in period 2. Their solution consequently fails to appreciate any reputation effects.

Polinsky and Rubinfeld (1991) study a setting with perfect monitoring. They assume that an individual's gain from committing some crime has two components: a socially acceptable gain and an illicit gain. The latter is a fixed trait. By contrast, an individual's acceptable gain is drawn from some distribution at the start of each period. Both components are private information. The planner maximizes aggregate acceptable gains minus harms stemming from criminal activities by choosing fines for first and second offenses. Since some crimes are socially efficient, the planner never opts for full deterrence. Individuals who commit crimes in the first period are likely to enjoy high illicit gains, especially if the fine for first offenses is low. This allows the planner to sort agents by 'illicit type'. Using higher fines for second offenses reduces underdeterrence vis-à-vis low uniform fines and reduces overdeterrence vis-à-vis high uniform fines, making such graduated fines socially optimal for some parameter values.[20]

Unlike Polinsky and Rubinfeld (1991), Polinsky and Shavell (1998) only consider acceptable gains in their two-period model with perfect monitoring. Polinsky and Shavell's planner has to expend resources to apprehend offenders and punishments cannot exceed some upper bound. Because administering punishments itself is costless, the planner uses this maximal punishment should using a uniform punishment be optimal. Since employing graduated punishments creates a reputation effect (a difference in tomorrow's punishments for first-time and repeat offenders makes agents more reluctant to commit a crime today), it can be optimal to set the punishment for first-time offenders in the second period below the maximal punishment. This reputation effect increases crime deterrence in period one, but reduces deterrence in period two. Whether the positive period-one effect outweighs the negative period-two effect depends on the distribution from which acceptable gains are

[20]If acceptable gains were fixed and illicit gains were drawn at the start of each period, then it can be optimal to use *lower* fines for second offenses.

drawn. In contrast to our reputation effect, Polinsky and Shavell's reputation effect has no impact on the punishment that prevails in period one.

In Rubinstein (1980) an agent's income should he abide the law is stochastic. Because the probability that the agent is caught when committing a crime is less than one, his income from criminal activities is also stochastic. Whether a uniform punishment scheme (in which punishments for first-time offenders and repeat offenders equal the maximal punishment) or a graduated punishment scheme is best at minimizing the number of offenses depends on the agent's risk attitude.

Warnings play a prominent role in Harrington (1988). Harrington, studying the enforcement of compliance with environmental regulations, shows that a planner who knows each firm's cost of compliance can achieve a higher compliance rate (compared to a system with a uniform punishment) by resorting to a system in which firms with relatively good compliance records are merely warned. Just like in our model, firms do not want to lose their good reputation, i.e. move to a high punishment regime. Yet, since Harrington (1988) assumes perfect monitoring, this result hinges on the presence of an upper bound to punishments. In a more recent paper, Rousseau (2009) argues that the use of warnings reduces the number of erroneous convictions and at the same time mitigates overcompliance to regulations by low types. Importantly, Rousseau assumes that the structure of punishments is exogenously given and that the planner can only choose between administering the appropriate punishment and warning the alleged violator.

Landsberger and Meilijson (1982) study how tax evasion is best combatted in a dynamic setting with an exogenously given penalty system and a homogeneous population. The tax authority is resource-constrained and can hence only audit a fraction of the population. Landsberger and Meilijson show that if the tax authority is sufficiently resource-constrained, then tax revenues are higher (compared to a uniform probability of being audited) if those who have been caught evading taxes in the previous period are audited with a higher probability than those who have not been caught evading taxes in the previous period.

Our approach is related to the model developed by Abreu et al. (2005). They study ongoing relationships between two players in which one player is tempted to depart from jointly efficient behaviour. How tempted that player is is private information. The other player receives signals regarding the tempted player's behaviour and can administer punishments to that player. In equilibrium punishments can go in either direction after perceived bad behaviour. The sign of the change in punishment depends crucially on the distribution from which the level of temptation is drawn. Although Abreu et al. (2005) stress that both asymmetric information and imperfect monitoring are a prerequisite for graduated punishments to occur, the setting they consider differs considerably from ours. They investigate a one-sided prisoner's dilemma with players who try to maximize their own payoff. In our public good game only the agents are selfish, the planner is benevolent. More importantly, the player who is tempted to depart from jointly efficient behaviour is infinitely impatient. As a consequence, reputation effects do not play a role in Abreu et al. (2005).

7 Concluding Remarks

We have investigated the optimal punishment scheme a social planner uses when confronted with a repeated public good problem. Because monitoring is imperfect and administering punishments is costly, a uniform punishment is often suboptimal. To alleviate the detrimental effects on welfare of monitoring mistakes and costly punishments, the planner employs a punishment scheme featuring graduated punishments: repeat offenders are punished harsher than first-time offenders. Such a punishment scheme allows the planner to (imperfectly) sort agents by cost type, enabling her to tailor future punishments to type. Moreover, because agents fear becoming branded as shirkers, i.e getting a bad reputation, the planner can allow herself to sanction first-time offenders very mildly. In fact, merely warning first-time offenders often suffices.

Obviously, one can envision more elaborate punishment schemes. For instance, in most judiciary systems the punishment a convicted criminal receives does not simply depend on whether this person already has a criminal record, but also on the precise content of such a record. Furthermore, we have only looked at the stationary equilibria of the infinite-horizon setting. We have consequently left an important question unanswered: under what conditions do groups or societies reach steady states in which graduated punishments are employed? Analysis of the short run-properties of a repeated game akin to the one discussed in Section 5 could help answering this question. These issues might prove fruitful avenues for future research.

Appendix

Details regarding Condition 1

Suppose that $\rho = 0$. To induce agents to contribute the planner has to set a punishment f such that $\gamma + \epsilon_{II} f \leq (1 - \epsilon_I) f$ and hence the planner opts for $f^* = \frac{\gamma}{1-\epsilon_I-\epsilon_{II}}$. The associated welfare reads

$$W(f^*) = 1 - \gamma - \tfrac{\epsilon_{II}}{1-\epsilon_I-\epsilon_{II}}(c+m)\gamma,$$

which is positive if $1 - \gamma > \frac{\epsilon_{II}}{1-\epsilon_I-\epsilon_{II}}(c+m)\gamma$ holds.

Proof of Proposition 1

Welfare with the low punishment equals

$$W(\tfrac{\gamma-\alpha}{\phi}) = \rho(1-\gamma+\alpha) - \rho\tfrac{\epsilon_{II}}{\phi}(c+m)(\gamma-\alpha) - (1-\rho)\tfrac{1-\epsilon_I}{\phi}c(\gamma-\alpha), \tag{A.1}$$

where we used the fact that $\delta_L = 1$ and $\delta_H = 0$ if $\phi f_0 = \gamma - \alpha$.

If the planner uses the high punishment, then $\delta_L = \delta_H = 1$ and hence welfare becomes

$$W(\tfrac{\gamma}{\phi}) = \rho(1-\gamma+\alpha) + (1-\rho)(1-\gamma) - \tfrac{\epsilon_{II}}{\phi}(c+m)\gamma. \tag{A.2}$$

The difference in welfare $\Delta = \Delta(\rho) := W(\frac{\gamma}{\phi}) - W(\frac{\gamma-\alpha}{\phi})$ between the two options reads

$$\begin{aligned}\Delta =& (1-\rho)(1-\gamma) - \tfrac{\epsilon_{II}}{\phi}(c+m)\gamma + \rho\tfrac{\epsilon_{II}}{\phi}(c+m)(\gamma-\alpha) + (1-\rho)\tfrac{1-\epsilon_I}{\phi}c(\gamma-\alpha)\\ =& (1-\rho)(1-\gamma) - \tfrac{\epsilon_{II}}{\phi}(c+m)\alpha + (1-\rho)c(\gamma-\alpha) - (1-\rho)\tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha).\end{aligned}$$

Solving $\Delta(\rho) = 0$ yields $\rho = \bar{\rho}$. Note that

$$\begin{aligned}\Delta'(\rho) =& -(1-\gamma) - c(\gamma-\alpha) + \tfrac{\epsilon_{II}}{\phi} m(\gamma-\alpha) < -\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma - c(\gamma-\alpha) + \tfrac{\epsilon_{II}}{\phi} m(\gamma-\alpha) \\ =& -\tfrac{\epsilon_{II}}{\phi} c(\gamma+\alpha) - c(\gamma-\alpha) < 0,\end{aligned}$$

where we used Condition 1 to establish the first inequality. Furthermore, $\Delta(1) = -\frac{\epsilon_{II}}{\phi}(c+m)\alpha < 0$ and

$$\Delta(0) = 1-\gamma + c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi} c\alpha - \tfrac{\epsilon_{II}}{\phi} m\gamma > 1-\gamma + c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}(c+m)\gamma > c(\gamma-\alpha),$$

where the first inequality follows from the fact that $\gamma > \alpha$ and the second one from Condition 1. The above observations imply that the planner opts for the punishment $\frac{\gamma-\alpha}{\phi}$ if $\rho \leq \bar{\rho}$ whereas she opts for the punishment $\frac{\gamma}{\phi}$ if $\rho > \bar{\rho}$ and that $\bar{\rho} \in (0,1)$. ■

Proof of Proposition 2

The total welfare $\mathcal{W}(f_0^*)$ generated if the planner uses the single punishment f_0^* is

$$\mathcal{W}(f_0^*) = \begin{cases} 2W(\frac{\gamma}{\phi}) & \text{if } \rho \leq \bar{\rho} \\ 2W(\frac{\gamma-\alpha}{\phi}) & \text{if } \rho > \bar{\rho}, \end{cases}$$

where $W(\frac{\gamma}{\phi})$ and $W(\frac{\gamma-\alpha}{\gamma})$ can be found in (A.2) respectively (A.1). We have to compare $\mathcal{W}(f_0^*)$ with $\mathcal{W}(\boldsymbol{f}^*) = W_1(\boldsymbol{f}^*) + W_2(\boldsymbol{f}^*)$.

In period 2 all low-cost types as well as those high-cost types who were caught shirking in period 1, i.e. a fraction $1-\epsilon_I$ of the high-cost types, contribute. This yields, after taking into account agents' costs of contributing, an aggregate payoff of

$$\rho(1-\gamma+\alpha) + (1-\rho)(1-\epsilon_I)(1-\gamma).$$

We have to deduct the social costs of administering punishments from this figure. These costs amount to

$$F_2(\boldsymbol{f}^*) = \rho\epsilon_{II}^2(c+m)\hat{f}_2^* + \rho(1-\epsilon_{II})\epsilon_{II}(c+m)f_2^* + (1-\rho)(1-\epsilon_I)\epsilon_{II}(c+m)\hat{f}_2^* + (1-\rho)\epsilon_I(1-\epsilon_I)cf_2^*.$$

So, $W_2(\boldsymbol{f}^*) = \rho(1-\gamma+\alpha) + (1-\rho)(1-\epsilon_I)(1-\gamma) - F_2(\boldsymbol{f}^*)$.

In period 1 only the low-cost types contribute, yielding an aggregate payoff of $\rho(1-\gamma+\alpha)$. The social costs of administering punishments in period 1 read

$$F_1(f_1^*) = \rho\epsilon_{II}(c+m)f_1^* + (1-\rho)(1-\epsilon_I)cf_1^*.$$

Welfare in period 1 thus reads $W_1(\boldsymbol{f}^*) = \rho(1-\gamma+\alpha) - F_1(f_1^*)$ and total welfare equals

$$\mathcal{W}(\boldsymbol{f}^*) = 2\rho(1-\gamma+\alpha) + (1-\rho)(1-\epsilon_I)(1-\gamma) - F_1(f_1^*) - F_2(\boldsymbol{f}^*). \tag{A.3}$$

The punishment f_1^* is either $\frac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}$ or 0 whereas f_0^* is either $\frac{\gamma}{\phi}$ or $\frac{\gamma-\alpha}{\phi}$. Let us now investigate the four parameter regions leading to these four combinations of f_1^* and f_0^*:

- $\gamma - \alpha - \epsilon_{II}\alpha \geq 0$ and $\rho \leq \bar{\rho}$: In this case $f_1^* = \frac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}$ and $f_0^* = \frac{\gamma}{\phi}$. Straightforward calculations reveal that

$$\begin{aligned} F_1(f_1^*) + F_2(\boldsymbol{f}^*) =& \big(2\rho\epsilon_{II} + (1-\rho)(1-\epsilon_I)\epsilon_{II}\big)(c+m)\tfrac{\gamma-\alpha}{\phi} + (1-\rho)(1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ &+ (1-\rho)(1-\epsilon_I)(1+\epsilon_I)c\tfrac{\gamma-\alpha}{\phi}. \end{aligned} \tag{A.4}$$

Subtracting $\mathcal{W}(f_0^*) = 2W(\frac{\gamma}{\phi})$ from $\mathcal{W}(\boldsymbol{f}^*)$ now yields

$$\begin{aligned} \Delta :=& \mathcal{W}(\boldsymbol{f}^*) - \mathcal{W}(f_0^*) = (1-\rho)(1-\epsilon_I)(1-\gamma) - 2(1-\rho)(1-\gamma) \\ &- \big(2\rho\epsilon_{II} + (1-\rho)(1-\epsilon_I)\epsilon_{II}\big)(c+m)\tfrac{\gamma-\alpha}{\phi} - (1-\rho)(1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ &- (1-\rho)(1-\epsilon_I)(1+\epsilon_I)c\tfrac{\gamma-\alpha}{\phi} + 2\epsilon_{II}(c+m)\tfrac{\gamma}{\phi} \\ =& -(1+\epsilon_I)(1-\rho)\big(1-\gamma+c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\big) \\ &+ 2\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} - (1-\rho)(1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi}, \end{aligned}$$

where we used $\phi = 1 - \epsilon_I - \epsilon_{II}$ to establish the second equality. Note that $\Delta = \Delta(\rho)$ is strictly increasing in ρ. Setting $\rho = 0$ gives us

$$\begin{aligned} \Delta(0) =& -(1+\epsilon_I)\big(1-\gamma+c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\big) + 2\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} - (1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ <& -(1+\epsilon_I)\big(\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma + c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\big) + 2\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} - (1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ =& -(1+\epsilon_I)\epsilon_{II}c\tfrac{\gamma}{\phi} - c(\gamma-\alpha) + 2\epsilon_{II}c\tfrac{\alpha}{\phi} \\ \leq& -\big((1+\epsilon_I)\epsilon_{II} + \phi\big)c(1+\epsilon_{II})\tfrac{\alpha}{\phi} + \big(\phi + 2\epsilon_{II}\big)c\tfrac{\alpha}{\phi} = -\epsilon_I\epsilon_{II}^2\tfrac{\alpha}{\phi} < 0, \end{aligned}$$

where the first inequality follows from Condition 1 and the second inequality follows from the fact that $\gamma - \alpha - \epsilon_{II}\alpha \geq 0$. Furthermore:

$$\begin{aligned} \Delta(\bar{\rho}) =& -(1+\epsilon_I)\tfrac{\epsilon_{II}}{\phi}(c+m)\alpha + 2\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} - (1-\bar{\rho})(1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ =& (1-\epsilon_I)\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} - (1-\bar{\rho})(1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} > 0. \end{aligned}$$

So, $\mathcal{W}(\boldsymbol{f}^*) > \mathcal{W}(f_0^*)$ if $\rho \in (\ell_+, \bar{\rho}]$ for some lower bound $\ell_+ \in (0, \bar{\rho})$.

- $\gamma - \alpha - \epsilon_{II}\alpha \geq 0$ and $\rho > \bar{\rho}$: In this case $f_1^* = \frac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}$ and $f_0^* = \frac{\gamma-\alpha}{\phi}$. The social costs of administering punishments are again those given in (A.4). Subtracting $2W(\frac{\gamma-\alpha}{\phi})$ from $\mathcal{W}(\boldsymbol{f}^*)$ results in

$$\begin{aligned} \Delta =& 2\rho(1-\gamma+\alpha) + (1-\rho)(1-\epsilon_I)(1-\gamma) - 2\rho(1-\gamma+\alpha) \\ &- \big(2\rho\epsilon_{II} + (1-\rho)(1-\epsilon_I)\epsilon_{II}\big)(c+m)\tfrac{\gamma-\alpha}{\phi} - (1-\rho)(1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ &- (1-\rho)(1-\epsilon_I)(1+\epsilon_I)c\tfrac{\gamma-\alpha}{\phi} + 2\rho\epsilon_{II}(c+m)\tfrac{\gamma-\alpha}{\phi} + 2(1-\rho)(1-\epsilon_I)c\tfrac{\gamma-\alpha}{\phi} \\ =& (1-\epsilon_I)(1-\rho)\big(1-\gamma+c(\gamma-\alpha) - \epsilon_{II}m\tfrac{\gamma}{\phi}\big) > 0, \end{aligned}$$

where the inequality follows from Condition 1. We conclude that in this case the planner always opts for $\boldsymbol{f}^*$.

- $\gamma-\alpha-\epsilon_{II}\alpha < 0$ and $\rho \le \bar{\rho}$: In this case $f_1^* = 0$ (implying that $F_1(f_1^*) = 0$) and $f_0^* = \frac{\gamma}{\phi}$. Denote the total welfare generated if the planner uses $\boldsymbol{f}^*$ when $\gamma - \alpha - \epsilon_{II}\alpha \ge 0$ by $\overline{\mathcal{W}(\boldsymbol{f}^*)}$. Then the welfare difference when $\gamma - \alpha - \epsilon_{II}\alpha < 0$ can be written as

  $$\Delta = \overline{\mathcal{W}(\boldsymbol{f}^*)} - \mathcal{W}(f_0^*) + F_1(\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}),$$

  where $F_1(\frac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}) < 0$. If $\rho = 0$, then the welfare difference becomes

  $$\begin{aligned}\Delta(0) =& -(1+\epsilon_I)\big(1-\gamma+c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\big) + 2\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} - (1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ &+ (1-\epsilon_I)c\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} \\ <& -(1+\epsilon_I)\big(\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma + c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\big) + 2\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} - (1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ &+ (1-\epsilon_I)c\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} = -\epsilon_I\epsilon_{II}c\tfrac{\gamma-\alpha}{\phi} < 0.\end{aligned}$$

  Evaluating the difference Δ at $\rho = \bar{\rho}$ gives us

  $$\begin{aligned}\Delta(\bar{\rho}) =& (1-\epsilon_I)\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} - (1-\bar{\rho})(1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ &+ \bar{\rho}\epsilon_{II}(c+m)\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} + (1-\bar{\rho})(1-\epsilon_I)c\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} \\ =& \bar{\rho}(1-\epsilon_I)\epsilon_{II}(c+m)\tfrac{\alpha}{\phi} + \bar{\rho}\epsilon_{II}(c+m)\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} + (1-\bar{\rho})(1-\epsilon_I)c\tfrac{\gamma-\alpha}{\phi} \\ =& \bar{\rho}\epsilon_{II}(c+m)\alpha + \bar{\rho}\epsilon_{II}(c+m)\tfrac{\gamma-\alpha}{\phi} + (1-\bar{\rho})(1-\epsilon_I)c\tfrac{\gamma-\alpha}{\phi} > 0,\end{aligned}$$

  where we used $\phi = 1 - \epsilon_I - \epsilon_{II}$ to establish the last equality. Differentiating Δ with respect to ρ yields

  $$\begin{aligned}\Delta'(\rho) =& (1+\epsilon_I)\big(1-\gamma+c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\big) + (1-\epsilon_I)\epsilon_{II}m\tfrac{\alpha}{\phi} \\ &+ \epsilon_{II}(c+m)\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} - (1-\epsilon_I)c\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} \\ =& (1+\epsilon_I)\big(1-\gamma+c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\big) + \epsilon_{II}(c+m)\alpha + \epsilon_{II}m\tfrac{\gamma-\alpha}{\phi} - c(\gamma-\alpha) \\ =& (1+\epsilon_I)(1-\gamma) + \epsilon_I c(\gamma-\alpha) - \epsilon_I\epsilon_{II}m\tfrac{\gamma-\alpha}{\phi} + \epsilon_{II}(c+m)\alpha > 0,\end{aligned}$$

  where the inequality is a consequence of Condition 1. We conclude that $\mathcal{W}(\boldsymbol{f}^*) > \mathcal{W}(f_0^*)$ if $\rho \in (\ell_0, \bar{\rho}]$ for some lower bound $\ell_0 \in (0, \bar{\rho})$.

- $\gamma - \alpha - \epsilon_{II}\alpha < 0$ and $\rho > \bar{\rho}$: The welfare difference now reads

  $$\begin{aligned}\Delta =& \overline{\mathcal{W}(\boldsymbol{f}^*)} - \mathcal{W}(f_0^*) + F_1(\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}) = (1-\epsilon_I)(1-\rho)\big(1-\gamma+c(\gamma-\alpha) - \epsilon_{II}m\tfrac{\gamma}{\phi}\big) \\ &+ \rho\epsilon_{II}(c+m)\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} + (1-\rho)(1-\epsilon_I)c\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}.\end{aligned}$$

  Differentiating Δ with respect to ρ results in

  $$\Delta'(\rho) = -(1-\epsilon_I)\big(1-\gamma+c(\gamma-\alpha) - \epsilon_{II}m\tfrac{\gamma}{\phi} + c\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}\big) + \epsilon_{II}(c+m)\tfrac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi}.$$

  We want to show that $\Delta'(\rho) < 0$. Because $\epsilon_{II}(c+m)\frac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} < 0$, it suffices to prove that $\chi := 1-\gamma+c(\gamma-\alpha) - \epsilon_{II}m\frac{\gamma}{\phi} + c\frac{\gamma-\alpha-\epsilon_{II}}{\phi} > 0$. Using the fact that $\gamma > \alpha$ and Condition 1 one obtains:

  $$\chi > 1-\gamma+c(\gamma-\alpha) + c\tfrac{\gamma-\alpha}{\phi} - \epsilon_{II}(c+m)\tfrac{\gamma}{\phi} > c(\gamma-\alpha) + c\tfrac{\gamma-\alpha}{\phi} > 0.$$

So, $\Delta'(\rho) < 0$. The analysis of the case $\gamma - \alpha - \epsilon_{II}\alpha < 0$ and $\rho \leq \bar{\rho}$ revealed that $\Delta(\bar{\rho}) > 0$. Since $\lim_{\rho \downarrow \bar{\rho}} \Delta(\rho) = \Delta(\bar{\rho}) > 0$ and $\Delta'(\rho) < 0$ for $\rho \in (\bar{\rho}, 1]$, we conclude that $\mathcal{W}(\boldsymbol{f}^*) > \mathcal{W}(f_0^*)$ if $\rho \in [\bar{\rho}, u)$ for some upper bound $u \in (\bar{\rho}, 1]$. Because $\Delta(1) = \epsilon_{II}(c+m)\frac{\gamma-\alpha-\epsilon_{II}\alpha}{\phi} < 0$, u is less than 1.

The analysis of the four cases reveals that the planner maximizes total welfare by using the menu $\boldsymbol{f}^*$ if $\rho \in (\check{\rho}, \hat{\rho})$, where $\check{\rho}$ is either ℓ_+ or ℓ_0 and $\hat{\rho}$ is either 1 or u. ■

Details regarding (15)

If $\delta_L = 1$, then (10) becomes

$$C_L = \gamma - \alpha + \epsilon_{II}(f + \beta\hat{C}_L) + (1-\epsilon_{II})\beta C_L \Leftrightarrow C_L = \frac{\gamma - \alpha + \epsilon_{II} f + \epsilon_{II}\beta\hat{C}_L}{1-\beta(1-\epsilon_{II})},$$

from which one infers using (14) that

$$\hat{C}_L - C_L = \frac{(1-\beta)\hat{C}_L - (\gamma-\alpha) - \epsilon_{II} f}{1-\beta(1-\epsilon_{II})} = \frac{\epsilon_{II}(\frac{\gamma}{\phi} - f)}{1-\beta(1-\epsilon_{II})}.$$

Consequently:

$$\gamma - \alpha \leq \phi f + \phi\beta(\hat{C}_L - C_L) \Leftrightarrow \phi f \geq (\gamma - \alpha) - \beta\frac{\epsilon_{II}(\gamma - \phi f)}{1-\beta(1-\epsilon_{II})} \Leftrightarrow \phi f \geq \gamma - \alpha - \tfrac{\beta}{1-\beta}\epsilon_{II}\alpha.$$

Substituting $\delta_H = 0$ in (12) yields

$$C_H = (1-\epsilon_I)(f + \beta\hat{C}_H) + \epsilon_I\beta C_H \Leftrightarrow C_H = \frac{(1-\epsilon_I)(f+\beta\hat{C}_H)}{1-\beta\epsilon_I}.$$

Combining the last equality with (14) results in

$$\hat{C}_H - C_H = \frac{(1-\beta)\hat{C}_H - (1-\epsilon_I)f}{1-\beta\epsilon_I} = \frac{(1-\epsilon_I)(\frac{\gamma}{\phi} - f)}{1-\beta\epsilon_I}.$$

Therefore:

$$\gamma > \phi f + \phi\beta(\hat{C}_H - C_H) \Leftrightarrow \gamma > \phi f + \beta\frac{(1-\epsilon_I)(\gamma - \phi f)}{1-\beta\epsilon_I} \Leftrightarrow \phi f < \gamma.$$

Proof of Proposition 3

If the planner opts for graduated punishment, then she sets $\phi f^* = \max\{\gamma - \alpha - \frac{\beta}{1-\beta}\epsilon_{II}\alpha, 0\}$. We first derive the associated per-period welfare without the social costs of administering punishments. In each period all low-cost types (a fraction ρ of the population) as well as the high-cost types in $\hat{q}$ (a fraction $(1-q) - (\rho - \mu)$ of the population) contribute. Using Lemma 1 one obtains

$$(1-q) - (\rho - \mu) = (1-\rho)\frac{\beta(1-\epsilon_I)}{1-\beta\epsilon_I}.$$

The per-period welfare without the social costs of administering punishments hence equals

$$\Psi = \rho(1-\gamma+\alpha) + (1-\rho)\frac{\beta(1-\epsilon_I)}{1-\beta\epsilon_I}(1-\gamma).$$

We next derive the social costs of administering punishments $F_\infty = F_\infty(f^*, \hat{f}^*)$. In each period a fraction ϵ_{II} of the low-cost types in q and a fraction $1-\epsilon_I$ of the high-cost types in q receive the low punishment f^*. Moreover, a fraction ϵ_{II} of the agents in $\hat{q}$ receive the high punishment $\hat{f}^*$. Only the high-cost types in q are punished rightfully. So:

$$\begin{aligned}
F_\infty =& \mu\epsilon_{II}(c+m)f^* + (q-\mu)(1-\epsilon_I)cf^* + (1-q)\epsilon_{II}(c+m)\hat{f}^* \\
=& \frac{\rho(1-\beta)}{1-\eta}\epsilon_{II}(c+m)f^* + \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}(1-\epsilon_I)cf^* \\
&+ \frac{\beta(1-\epsilon_I)(1-\eta) - \rho\beta(1-\beta)\phi}{(1-\beta\epsilon_I)(1-\eta)}\epsilon_{II}(c+m)\hat{f}^* \\
=& \left(1 - \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}\right)\epsilon_{II}(c+m)\hat{f}^* - \frac{\rho(1-\beta)}{1-\eta}\epsilon_{II}(c+m)D \\
&+ \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}(1-\epsilon_I)c(\hat{f}^* - D) \\
=& \left(\epsilon_{II} + \frac{(1-\rho)(1-\beta)\phi}{1-\beta\epsilon_I}\right)c\hat{f}^* + \left(\epsilon_{II} - \frac{(1-\rho)(1-\beta)\epsilon_{II}}{1-\beta\epsilon_I}\right)m\hat{f}^* \\
&- \left(\frac{(1-\beta)\epsilon_{II}}{1-\eta} - \frac{(1-\rho)(1-\beta)\epsilon_{II}}{1-\eta}\right)mD - \left(\frac{(1-\beta)\epsilon_{II}}{1-\eta} + \frac{(1-\rho)(1-\beta)^2\phi}{(1-\eta)(1-\beta\epsilon_I)}\right)cD,
\end{aligned}$$

where the second equality follows from Lemma 1 and $D := \hat{f}^* - f^*$ is either $\frac{\alpha + \frac{\beta}{1-\beta}\epsilon_{II}\alpha}{\phi} = \frac{1-\eta}{(1-\beta)\phi}\alpha$ or 0. Of course, $\hat{f}^* = \frac{\gamma}{\phi}$.

The per-period welfare if graduated punishments are used is $\mathcal{W}_\infty = \Psi - F_\infty$. We have to compare this figure with $W(f_0^*)$, the per-period welfare if the single punishment f_0^* is used. Let us now analyze the difference $\Delta_\infty = \Delta_\infty(\rho) := \mathcal{W}_\infty - W(f_0^*)$ for the four cases that require attention:

- $\gamma - \alpha - \frac{\beta}{1-\beta}\epsilon_{II}\alpha \geq 0$, $\rho \leq \bar{\rho}$: In this case $D = \frac{1-\eta}{(1-\beta)\phi}\alpha$ and $W(f_0^*) = W(\frac{\gamma}{\phi})$. Hence:

$$\begin{aligned}
\Delta_\infty =& -\frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}(1-\gamma) - \epsilon_{II}c\tfrac{\gamma}{\phi} - \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}c\gamma - \epsilon_{II}m\tfrac{\gamma}{\phi} \\
&+ \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}\epsilon_{II}m\tfrac{\gamma}{\phi} - (1-\rho)\epsilon_{II}m\tfrac{\alpha}{\phi} + \epsilon_{II}m\tfrac{\alpha}{\phi} + \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}c\alpha \\
&+ \epsilon_{II}c\tfrac{\alpha}{\phi} + \epsilon_{II}(c+m)\tfrac{\gamma}{\phi} \\
=& -\frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}\left(1-\gamma+c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\right) \\
&- \frac{(1-\rho)\beta(1-\epsilon_I)}{1-\beta\epsilon_I}m\tfrac{\epsilon_{II}}{\phi}\alpha + (c+m)\tfrac{\epsilon_{II}}{\phi}\alpha.
\end{aligned}$$

Observe that Δ_∞ is increasing in ρ. Furthermore:

$$\begin{aligned}\Delta_\infty(\bar\rho) =& -\frac{1-\beta}{1-\beta\epsilon_I}(c+m)\tfrac{\epsilon_{II}}{\phi}\alpha - \frac{(1-\bar\rho)\beta(1-\epsilon_I)}{1-\beta\epsilon_I}m\tfrac{\epsilon_{II}}{\phi}\alpha + (c+m)\tfrac{\epsilon_{II}}{\phi}\alpha\\ =&\frac{\beta(1-\epsilon_I)}{1-\beta\epsilon_I}\left((c+m)\tfrac{\epsilon_{II}}{\phi}\alpha - (1-\bar\rho)m\tfrac{\epsilon_{II}}{\phi}\alpha\right) > 0.\end{aligned}$$

We conclude that $\mathcal{W}_\infty > W(f_0^*)$ if $\rho \in (\ell_+, \bar\rho]$ for some lower bound $\ell_+ \in [0, \bar\rho)$. One easily verifies that $\ell_+ = 0$ for β sufficiently close to 1.

- $\gamma - \alpha - \frac{\beta}{1-\beta}\epsilon_{II}\alpha \geq 0$, $\rho > \bar\rho$: We now have $D = \frac{1-\eta}{(1-\beta)\phi}\alpha$ and $W(f_0^*) = W(\frac{\gamma-\alpha}{\phi})$. Consequently:

$$\begin{aligned}\Delta_\infty =& \left(1 - \frac{1-\beta}{1-\beta\epsilon_I}\right)(1-\rho)(1-\gamma) - \epsilon_{II}c\tfrac{\gamma}{\phi} - \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}c\gamma - \epsilon_{II}m\tfrac{\gamma}{\phi}\\ &+ \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}\epsilon_{II}m\tfrac{\gamma}{\phi} - (1-\rho)\epsilon_{II}m\tfrac{\alpha}{\phi} + \epsilon_{II}m\tfrac{\alpha}{\phi} + \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}c\alpha\\ &+ \epsilon_{II}c\tfrac{\alpha}{\gamma} + \rho\tfrac{\epsilon_{II}}{\phi}(c+m)(\gamma-\alpha) + (1-\rho)\tfrac{1-\epsilon_I}{\phi}c(\gamma-\alpha)\\ =& \left(1 - \frac{1-\beta}{1-\beta\epsilon_I}\right)(1-\rho)(1-\gamma) - \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}c(\gamma-\alpha)\\ &+ \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}\tfrac{\epsilon_{II}}{\phi}m\gamma - (1-\rho)\tfrac{\epsilon_{II}}{\phi}m\gamma - (1-\rho)c(\gamma-\alpha)\\ =&\frac{(1-\rho)\beta(1-\epsilon_I)}{1-\beta\epsilon_I}\left(1-\gamma+c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m\gamma\right) > 0,\end{aligned}$$

where the inequality follows from Condition 1. So, $\mathcal{W}_\infty > W(f_0^*)$ for all $\rho \in (\bar\rho, 1]$.

- $\gamma - \alpha - \frac{\beta}{1-\beta}\epsilon_{II}\alpha < 0$, $\rho \leq \bar\rho$: In this case $D = \frac{\gamma}{\phi}$ and $W(f_0^*) = W(\frac{\gamma}{\phi})$. Therefore:

$$\begin{aligned}\Delta_\infty =& -\frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}(1-\gamma) - \epsilon_{II}c\tfrac{\gamma}{\phi} - \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}c\gamma - \epsilon_{II}m\tfrac{\gamma}{\phi}\\ &+ \frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}\epsilon_{II}m\tfrac{\gamma}{\phi} - \frac{(1-\rho)(1-\beta)}{1-\eta}\epsilon_{II}m\tfrac{\gamma}{\phi} + \frac{1-\beta}{1-\eta}\epsilon_{II}m\tfrac{\gamma}{\phi}\\ &+ \frac{(1-\rho)(1-\beta)^2}{(1-\eta)(1-\beta\epsilon_I)}c\gamma + \frac{1-\beta}{1-\eta}\epsilon_{II}c\tfrac{\gamma}{\phi} + \epsilon_{II}(c+m)\tfrac{\gamma}{\phi}\\ =& -\frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}(1-\gamma) - \frac{(1-\rho)\beta(1-\beta)\epsilon_{II}}{(1-\eta)(1-\beta\epsilon_I)}c\gamma\\ &- \frac{(1-\rho)\beta(1-\beta)\epsilon_{II}}{(1-\eta)(1-\beta\epsilon_I)}m\gamma + \frac{1-\beta}{1-\eta}\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma\\ =&\frac{1-\beta}{1-\beta\epsilon_I}\left(-(1-\rho)(1-\gamma) - \frac{(1-\rho)\beta\epsilon_{II}}{1-\eta}(c+m)\gamma + \frac{1-\beta\epsilon_I}{1-\eta}\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma\right).\end{aligned}$$

Clearly, Δ_∞ is increasing in ρ. Using Condition 1 one infers that:

$$\begin{aligned}\Delta_\infty(0) =&\frac{1-\beta}{1-\beta\epsilon_I}\left(-(1-\gamma)-\frac{\beta\epsilon_{II}}{1-\eta}(c+m)\gamma+\frac{1-\beta\epsilon_I}{1-\eta}\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma\right)\\ =&\frac{1-\beta}{1-\beta\epsilon_I}\left(-(1-\gamma)+\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma\right)<0.\end{aligned}$$

We now prove that $\Delta_\infty(\bar\rho)>0$, i.e. that

$$\omega(\bar\rho):=-(1-\bar\rho)(1-\gamma)-\frac{(1-\bar\rho)\beta\epsilon_{II}}{1-\eta}(c+m)\gamma+\frac{1-\beta\epsilon_I}{1-\eta}\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma>0.$$

Since

$$\frac{\mathrm{d}(1-\bar\rho)}{\mathrm{d}\alpha}=\frac{\tfrac{\epsilon_{II}}{\phi}(c+m)(1-\gamma+c\gamma-\tfrac{\epsilon_{II}}{\phi}m\gamma)}{\left(1-\gamma+c(\gamma-\alpha)-\tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\right)^2}>0$$

and $\omega'(\bar\rho)<0$, we have that $\omega(\bar\rho)>\omega(\lim_{\alpha\to\gamma}\bar\rho)$. In other words, replacing α by its upper bound γ in the expression for $\bar\rho$ yields a lower bound for $\omega(\bar\rho)$. Using the fact that

$$\lim_{\alpha\to\gamma}(1-\bar\rho)=\frac{\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma}{1-\gamma}$$

one obtains

$$\begin{aligned}\omega(\bar\rho)>&-\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma-\frac{\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma}{1-\gamma}\times\frac{\beta\epsilon_{II}}{1-\eta}(c+m)\gamma+\frac{1-\beta\epsilon_I}{1-\eta}\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma\\ =&\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma\left(-1-\frac{\beta\epsilon_{II}}{(1-\eta)(1-\gamma)}(c+m)\gamma+\frac{1-\beta\epsilon_I}{1-\eta}\right)\\ >&\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma\left(-1-\frac{\beta\phi}{1-\eta}+\frac{1-\beta\epsilon_I}{1-\eta}\right)=0,\end{aligned}$$

where we used Condition 1 to establish the second inequality. We conclude that $\mathcal{W}_\infty>W(f_0^*)$ if $\rho\in(\ell_0,\bar\rho]$ for some lower bound $\ell_0\in(0,\bar\rho)$.

- $\gamma-\alpha-\frac{\beta}{1-\beta}\epsilon_{II}\alpha<0$, $\rho>\bar\rho$: We now have $D=\frac{\gamma}{\phi}$ and $W(f_0^*)=W(\frac{\gamma-\alpha}{\phi})$. Hence:

$$\begin{aligned}\Delta_\infty=&(1-\rho)(1-\gamma)-\frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}(1-\gamma)-\epsilon_{II}c\tfrac{\gamma}{\phi}-\frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}c\gamma-\epsilon_{II}m\tfrac{\gamma}{\phi}\\ &+\frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}\epsilon_{II}m\tfrac{\gamma}{\phi}-\frac{(1-\rho)(1-\beta)}{1-\eta}\epsilon_{II}m\tfrac{\gamma}{\phi}+\frac{1-\beta}{1-\eta}\epsilon_{II}m\tfrac{\gamma}{\phi}\\ &+\frac{(1-\rho)(1-\beta)^2}{(1-\eta)(1-\beta\epsilon_I)}c\gamma+\frac{(1-\beta)\epsilon_{II}}{1-\eta}c\tfrac{\gamma}{\phi}+\rho\tfrac{\epsilon_{II}}{\phi}(c+m)(\gamma-\alpha)+(1-\rho)\tfrac{1-\epsilon_I}{\phi}c(\gamma-\alpha)\\ =&(1-\rho)\left(1-\gamma+c(\gamma-\alpha)-\tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha)\right)-\frac{(1-\rho)(1-\beta)}{1-\beta\epsilon_I}(1-\gamma)\\ &-\tfrac{\epsilon_{II}}{\phi}(c+m)\alpha-\frac{(1-\rho)\beta(1-\beta)\epsilon_{II}}{(1-\eta)(1-\beta\epsilon_I)}(c+m)\gamma+\frac{(1-\beta)}{1-\eta}\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma.\end{aligned}$$

Because $\gamma < \alpha + \frac{\beta}{1-\beta}\epsilon_{II}\alpha = \frac{1-\eta}{1-\beta}\alpha$, we have that

$$\Delta_\infty(1) = \tfrac{\epsilon_{II}}{\phi}(c+m)\left(-\alpha + \frac{1-\beta}{1-\eta}\gamma\right) < 0.$$

Differentiating Δ_∞ with respect to ρ yields

$$\begin{aligned}
\Delta'_\infty(\rho) &= -\frac{\beta(1-\epsilon_I)}{1-\beta\epsilon_I}(1-\gamma) - c(\gamma-\alpha) + \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha) + \frac{\beta(1-\beta)\epsilon_{II}}{(1-\eta)(1-\beta\epsilon_I)}(c+m)\gamma \\
&< -\frac{\beta(1-\epsilon_I)}{1-\beta\epsilon_I}\tfrac{\epsilon_{II}}{\phi}(c+m)\gamma - c(\gamma-\alpha) + \tfrac{\epsilon_{II}}{\phi}m(\gamma-\alpha) + \frac{\beta(1-\beta)\epsilon_{II}}{(1-\eta)(1-\beta\epsilon_I)}(c+m)\gamma \\
&= -\frac{\beta\epsilon_{II}}{1-\eta}\tfrac{\epsilon_{II}}{\phi}c\gamma + \frac{1-\beta}{1-\eta}\tfrac{\epsilon_{II}}{\phi}m\gamma - c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m\alpha \\
&< -\frac{\beta\epsilon_{II}}{1-\eta}\tfrac{\epsilon_{II}}{\phi}c\gamma + \tfrac{\epsilon_{II}}{\phi}m\alpha - c(\gamma-\alpha) - \tfrac{\epsilon_{II}}{\phi}m\alpha = -\frac{\beta\epsilon_{II}}{1-\eta}\tfrac{\epsilon_{II}}{\phi}c\gamma - c(\gamma-\alpha) < 0,
\end{aligned}$$

where the first inequality follows from Condition 1 and the second one from the fact that $\gamma < \frac{1-\eta}{1-\beta}\alpha$. Because $\Delta_\infty(\bar{\rho}) > 0$ (see the previous bullet point) and Δ_∞ is continuous in ρ, we conclude that $\mathcal{W}_\infty > W(f_0*)$ if $\rho \in (\bar{\rho}, u)$ for some upper bound $u \in (\bar{\rho}, 1)$.

The analysis of the four cases reveals that the planner maximizes per-period welfare by using the pair of punishments $(f^*, \hat{f}^*)$ if $\rho \in (\check{\rho}_\infty, \hat{\rho}_\infty)$, where $\check{\rho}_\infty$ is either ℓ_+ or ℓ_0 and $\hat{\rho}_\infty$ is either 1 or u. ■

Acknowledgements

Financial support from the Belgian Federal Government through the IAP Project (contract 6/09) is gratefully acknowledged. This paper owes its existence to numerous discussions with Wouter Vergote. I would like to thank Remco van Eijkel, attendees at PET10, SMYE2011, EEA-ESEM 2011, and workshop participants at CEREC (Facultés universitaires Saint-Louis), at CORE (Université Catholique de Louvain), and at IEEF.

References

Abreu, D., D. Bernheim, and A. Dixit (2005): “Self-Enforcing Cooperation with Graduated Punishments,” Working paper, Princeton University.

Agrawal, A. (2003): “Sustainable Governance of Common-Pool Resources: Context, Methods, and Politics,” *Annual Review of Anthropology*, 32, 243–262.

Becker, G. S. (1968): "Crime and Punishment: An Economic Approach," *Journal of Political Economy*, 76, 169–217.

Chu, C. Y. C., S.-C. Hu, and T.-Y. Huang (2000): "Punishing Repeat Offenders more Severely," *International Review of Law and Economics*, 20, 127–140.

Dana, D. A. (2001): "Rethinking the Puzzle of Escalating Penalties for Repeat Offenders," *Yale Law Journal*, 110, 733–783.

Ellickson, R. (1991): *Order without Law: How Neighbors Settle Disputes*. Harvard University Press, Cambridge, MA.

Green, E., and R. Porter (1984): "Noncooperative Collusion under Imperfect Price Information," *Econometrica*, 52, 87–100.

Harrington, W. (1988): "Enforcement Leverage when Penalties are Restricted," *Journal of Public Economics*, 37, 29–53.

Landsberger, M., and I. Meilijson (1982): "Incentive Generating State Dependent Penalty System," *Journal of Public Economics*, 19, 333–352.

Miceli, T. J., and C. Bucci (2005): "A Simple Theory of Increasing Penalties for Repeat Offenders," *Review of Law and Economics*, 1.

Mungan, M. C. (2010): "Repeat Offenders: If they learn, we punish them more severely," *International Review of Law and Economics*, 30, 173–177.

Ostrom, E. (1990): *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, UK.

——— (2000): "Collective Action and the Evolution of Social Norms," *Journal of Economic Perspectives*, 14, 137–158.

Polinsky, A. M., and D. L. Rubinfeld (1991): "A Model of Optimal Fines for Repeat Offenders," *Journal of Public Economics*, 46, 291–306.

Polinsky, A. M., and S. Shavell (1998): "On Offense History and the Theory of Deterrence," *International Review of Law and Economics*, 18, 305–324.

Rousseau, S. (2009): "The Use of Warnings in the Presence of Errors," *International Review of Law and Economics*, 29, 191–201.

Rubinstein, A. (1979): "An Optimal Policy for Offenses that May Have Been Committed by Accident," in *Applied Game Theory*, ed. by S. Brams, A. Schotter, and G. Schwodiauer, pp. 406–413. Physica-Verlag, Würzburg.

——— (1980): "On an Anomaly of the Deterrent Effect of Punishment," *Economic Letters*, 6, 89–94.

Stigler, G. J. (1970): "The Optimum Enforcement of Laws," *Journal of Political Economy*, 78, 526–536.

Wade, R. (1994): *Village Republics: Economic Conditions for Collective Action in South India*. ICS Press, Oakland.