

The Evolution of Coercive Institutional Punishment

Alexander Isakov · David G. Rand

© Springer Science+Business Media, LLC 2011

Abstract Institutional punishment plays a central role in human societies. Yet research in evolutionary game theory has focused almost exclusively on peer punishment. Here we present a model for the evolution of institutional punishment. We consider a set of states (“kingdoms”), each consisting of a number of individuals (“subjects”) and a single leader (“king”). Subjects choose how much to pay to the king as tribute. The king chooses how much to punish his subjects based on their tribute payment level, in an effort to exact as much tribute as possible. We find the existence of both coercive Nash equilibria with punishment and high tribute payments, and noncoercive Nash equilibria with no punishment and no tribute payments. We also examine stochastic coevolutionary dynamics using agent-based simulations. We find that within a single state, the more intensely the king punishes, the more subjects evolve to pay in tribute. The king earns the most when both punishment is strong and subjects are accurate in their learning. When we consider co-evolution occurring at the level of the king as well as the subject, we see that kings evolve to punish heavily, as long as subjects are sufficiently accurate and frequent in their learning, and learn predominately from subjects in the same kingdom. If citizens have error-prone learning and/or are slow to update their strategies, however, selection leads to kings who punish little. Thus confusion is collectively beneficial for subjects. In sum, we show circumstances under which natural selection can favor the emergence of institutional punishment as a tool of coercion.

Keywords Institutional punishment · Evolutionary game theory · Coercion · Cooperation

A. Isakov

Department of Mathematics, Harvard University, Cambridge, MA 02138, USA
e-mail: isakov@fas.harvard.edu

D.G. Rand (✉)

Program for Evolutionary Dynamics, Department of Psychology and Berkman Center for Internet & Society, Harvard University, Cambridge MA 02138, USA
e-mail: drand@fas.harvard.edu

1 Introduction

In recent years, the human appetite for costly punishment has received a great deal of attention in the evolutionary game theory literature [2, 10, 11, 17, 18, 25, 27, 30, 35, 36]. Behavioral experiments demonstrate that many people are willing to pay to make others incur costs across a range of settings [1, 6–8, 13, 20, 26, 40]. Many scholars have argued that costly punishment has coevolved with cooperation [2, 10, 11, 17, 30], although that view has been questioned [4, 7, 25–27].

Almost all of this previous work has focused on peer punishment in symmetric games amongst equals. Yet, costly peer punishment (i.e., vigilante justice) is relatively rare in modern societies. Instead, the vast majority of punishment is carried out by the state rather than the individual. Despite its ubiquity, the evolution of this type of institutional punishment has received very little attention. One recent exception is [31], which studies “pool punishment,” where members of a public goods game group pay into a pool which automatically punishes low contributors.

Here we study a form of institutional punishment in which the state exists as a separate strategic entity from the individual. The interaction is asymmetric, with a strong state punishing a comparatively weak individual. Thus retaliation and “antisocial punishment,” which pose serious challenges for peer punishment [7, 14, 27], are not a significant issue in the context of institutionalized punishment. Secondly, the “2nd-order free rider” problem of peer punishment [2, 12, 21, 40], in which some players cooperate (to avoid being punished) but do not engage in punishment and thus outcompete the punishers, does not exist with institutional punishment: there is only a single entity (the state) with the ability to punish. Thirdly, this form of institutional punishment need not be linked with pro-sociality. A successful state needs only to make sure the individuals are supporting the state, not necessarily that the individuals are helping each other. The relationship between the state and the individual may look less like a cooperation game, and more like the relationship between dominant and submissive animals [5]. Here it can be self-interested/adaptive for the powerful party (the state) to use punishment to coerce the weaker party (the individual) into giving up resources.

In this paper, we introduce a simple model for the evolution of coercive institutional punishment. We consider a set of states (or “kingdoms”), each consisting of a number of individuals (“subjects”) ruled by a leader (“king”). Subjects choose how much to pay to the king as tribute. The king chooses how much to punish his subjects based on the amount of tribute paid. We study the coevolution of subjects and kings, and ask when natural selection leads to the emergence of institutional punishment by kings. In Sect. 2, we introduce the model. In Sect. 3, we present the results of our analysis. In Sect. 4, we conclude.

2 Model

We consider a world of N separate kingdoms, K_1, \dots, K_N . Each kingdom i is a collection of n subjects and one king. In every interaction period, each subject j in kingdom K_i produces 1 unit of income. That subject’s strategy $p_{i,j} \in [0, 1]$ defines the fraction of income she chooses to pay in tribute to the treasury of her king. The remaining $1 - p_{i,j}$ is kept by the subject.

Each king i then chooses how much to punish each subject j in his kingdom K_i , based on her tribute payment $p_{i,j}$. Kings earn the highest income by exacting as much tribute as possible from their subjects, and may choose to use punishment as a tool to enforce payment. A given king i ’s strategy s_i determines how strongly he punishes deviations from the

maximum tribute payment of $p = 1$. We use a linear punishment function, such that in kingdom K_i , the king pays a cost $s_i(1 - p_{i,j})$ to cause subject i to incur a cost $s_i(1 - p_{i,j})$. Thus, the more tribute p a given subject pays, the less she is punished; and the larger a given king's punishment strength s , the more punishment is inflicted for a given level of tribute $p < 1$.

In each interaction round, the payoff $\pi_{i,j}$ of subject j in kingdom K_i is therefore

$$\pi_{i,j} = (1 - s_i)(1 - p_{i,j}), \tag{1}$$

and the payoff Π_i of the king of kingdom K_i is

$$\Pi_i = \sum_{j=1}^n (p_{i,j} - s_i(1 - p_{i,j})) \tag{2}$$

Both subjects and kings update their respective strategies p and s through evolutionary processes. As we are considering the cultural evolution of institutions, this process represents social learning. In particular, we use the ‘‘pairwise comparison’’ process [34]. Here a learner and teacher are picked at random, and the learner adopts the teacher's strategy with a probability that is monotonically increasing in the teacher's payoff advantage. We begin by assuming that subjects learn only from other subjects within their kingdom, and later explore the effects of allowing learning from subjects in other kingdoms. The probability of subject j in kingdom K_i adopting the strategy of subject k in kingdom K_i is given by

$$P_s = \frac{1}{1 + \exp(-\beta(\pi_{i,k} - \pi_{i,j}))}, \tag{3}$$

where β represents the fidelity of learning among subjects (which we will call the level of subject ‘‘intelligence,’’ although this also involves other factors such as accuracy of information). The larger β is, the more likely it is that a subject will adopt the strategy of another with a higher payoff, and the less likely it is that she will adopt the strategy of another with a lower payoff. In the limit $\beta \rightarrow \infty$, we have deterministic dynamics, where the learner always copies if the teacher has a higher payoff, never copies if the teacher has a lower payoff, and copies half of the time if the payoffs are equal.

Kings learn from the kings of other kingdoms. The probability of the king of kingdom K_i adopting the strategy of the king of kingdom K_j is given by

$$P_k = \frac{1}{1 + \exp(-\beta_K(\Pi_j - \Pi_i))}, \tag{4}$$

where β_K represents the fidelity of learning among kings. Again, $\beta_K \rightarrow \infty$ gives a deterministic imitation dynamic.

Alternatively, with some small probability μ a ‘‘mutation’’ occurs, and the learner selects a new strategy at random. In the context of social learning, this represents experimentation or innovation. When a mutation occurs, subjects pick a new strategy p from a uniform distribution on $[0, 1]$. There is no a priori upper limit on kings' punishment strength s , but one must be assumed to implement mutation. In our main simulations, mutating kings pick a new strategy s from a uniform distribution on $[0, 3]$. We also replicate our analysis using a different upper bound on s and show the robustness of our results.

After every interaction round, a king learning round occurs with probability θ , and a subject learning round occurs with probability $1 - \theta$. In either case, one agent per kingdom

has the chance to update strategy. Thus, during king learning rounds, each king updates; and during subject learning rounds, one subject per kingdom updates.

Note that the subject and king update rules are invariant to the addition of a constant to all payoffs. Thus, for simplicity, we allow negative payoffs for both subjects and kings, and implicitly assume the addition of a sufficiently large constant to make the payoffs positive. Alternatively, the model could be interpreted as allowing kings to “go into debt” in order to punish, and allowing peasants to borrow to pay tribute.

We also note that this game setup is similar to a set of dictator games with punishment. Each subject decides how to divide her income between herself and the king (like in the dictator game), and the king can then respond to this decision using costly punishment. This also has some similarity to the ultimatum game, but our kings have a much richer range of punishment options than in the ultimatum game, where the only punishment is to destroy all of both players’ income. And unlike both the standard dictator or ultimatum setups, a single king plays simultaneously (and independently) with many subjects.

3 Results

We now consider the evolutionary dynamics of this system. In particular, we are interested in when natural selection will favor kings that use punishment to coerce subjects into paying tribute.

3.1 Nash Equilibrium Analysis

We begin with a static Nash equilibrium analysis. We note that no game occurs between kings, as one king’s strategy does not affect the payoff of other kings. Instead, constant selection occurs at the level of the king. Thus we perform a Nash equilibrium analysis of the game between subject and king within a single kingdom.

Examining the peasant payoff function given in (1), we see that peasant payoff is strictly decreasing in p if $s < 1$ (paying tribute is costly and thus always disadvantageous if punishment is weak); peasant payoff is unaffected by p in the knife-edge situation where $s = 1$ (the cost of paying tribute is exactly balanced out by the decrease in punishment); and peasant payoff is strictly increasing in p if $s > 1$ (punishment is sufficiently strong that the fine for paying less tribute is bigger than the amount saved by not paying). Therefore perfectly rational peasants will pay no tribute $p = 0$ when $s < 1$ and will pay the maximal amount of tribute $p = 1$ when $s > 1$.

Examining the king payoff function given in (2), we see that king payoff is strictly decreasing in s if $p < 1$ and king payoff is unaffected by s if $p = 1$ (punishing is always costly to the king unless the subjects pay the maximal tribute, in which case the king never actually has to punish).

Thus, we see that $p = s = 0$ is a Nash equilibrium. Because $s < 1$, subjects have no incentive to increase their tribute payments; and because $p < 1$, kings have no incentive to increase their punishment. This is a noncoercive equilibrium where subjects keep the fruits of their labor (i.e., pay 0 tribute), and the king collects no tribute and imposes no punishment.

We also see an infinite set of Nash equilibria that satisfy the condition $p = 1, s \geq 1$. Because $s \geq 1$, subjects have no incentive to decrease their tribute payments ($p = 1$ is payoff maximizing if $s > 1$, p does not effect subject payoff if $s = 1$); and because $p = 1$, the king never punishes, so s has no effect on his payoff, and he therefore has no incentive to deviate. Here the king punishes sufficiently strongly that subjects do best by giving up all of their

earnings in tribute to the king. As a result, the king never has to punish and reaps all that is produced in the kingdom.

The existence of this second set of equilibria suggests that evolution may lead to kings who use punishment to coerce their subjects. Although the coercive equilibrium is not subgame perfect (once a subject has made her tribute decision, the king has no incentive to actually punish), natural selection may favor kings who punish over kings who do not, if the punishment leads to high-tribute subjects replacing lower-tribute subjects.

We note that both the coercive and noncoercive equilibria are equally efficient, because the king never has to punish at the coercive equilibrium. However, even a small amount of error in execution, perception, or imitation can cause some subjects to make nonmaximal tribute payments, leading kings to punish. This will cause the coercive equilibrium to become less efficient than the noncoercive equilibrium. Thus the use of punishment will almost certainly lead to lower aggregate payoffs in this framework.

3.2 Evolutionary Dynamics: One Kingdom

We now examine the evolutionary dynamics of subject strategies within a single kingdom. Using subject populations of $n = 10$ and $n = 100$, we determine the evolutionary outcomes under different fixed strengths of institutionalized punishment. Specifically, we ask how the average subject's strategy p , subject's payoff π and king's payoff Π vary as functions of the strength of punishment s (which is exogenously imposed and fixed—no king learning occurs). We also examine the effect of subject intelligence β , conducting simulations using $\beta \in \{0.2, 0.8, 2, \infty\}$. In each simulation, we fix the subject mutation rate at $\mu_s = 0.05$, initialize each subject with a randomly drawn strategy from the interval $[0, 1]$, and simulate for 10^6 strategy updates. We then display the values of p , π , and Π averaged over the second half of the simulation, to avoid effects of the initial conditions.

The results for $n = 10$ and $n = 100$ are shown in Fig. 1. We notice several clear regularities.

First we consider subject outcomes. We see that in all cases, stronger punishment leads to larger tribute payments p (although for $\beta \rightarrow \infty$, the transition is abrupt at $s = 1$ rather than continuous). These larger tributes payments, together with more stringent punishment for $p < 1$, lead to lower subject payoffs π for higher strength of punishment s . Punishment compels the subjects to pay tribute, to the detriment of the subjects. High subject intelligence allows subjects to more fully coordinate on the equilibrium strategy. Therefore the effect of subject intelligence β on tribute payment p depends on the strength of punishment s : subject intelligence decreases tribute payments when punishment is weak and increases tribute payments when punishment is strong. When $s < 1$, the equilibrium strategy is for subjects to pay no tribute $p = 0$. Thus, the smarter the subjects are (the larger β is), the more quickly p decreases towards 0 as s decreases below 1. When $s > 1$, on the other hand, the equilibrium strategy is full tribute payment $p = 1$. Thus, the smarter the subjects are, the more quickly p increases toward 1 as s increases above 1. As higher intelligence allows the subjects to more quickly and fully adopt the payoff maximizing strategy, be it noncompliance or compliance, higher subject intelligence always increases subject payoffs.

Considering the king payoff Π gives a more nuanced picture. First let us consider the case of low subject intelligence, $\beta = 0.2$. Here the incentive created by punishment has little effect on subject behavior, because subject learning is largely dominated by random drift. As a result, it is not worth it for the king to invest any money in punishing. King payoff is thus monotonically decreasing in punishment strength. When subject intelligence is higher, however, the relationship between punishment strength and king payoff is nonmonotonic: the worst payoff strategy for kings is ambivalence, only punishing an intermediate

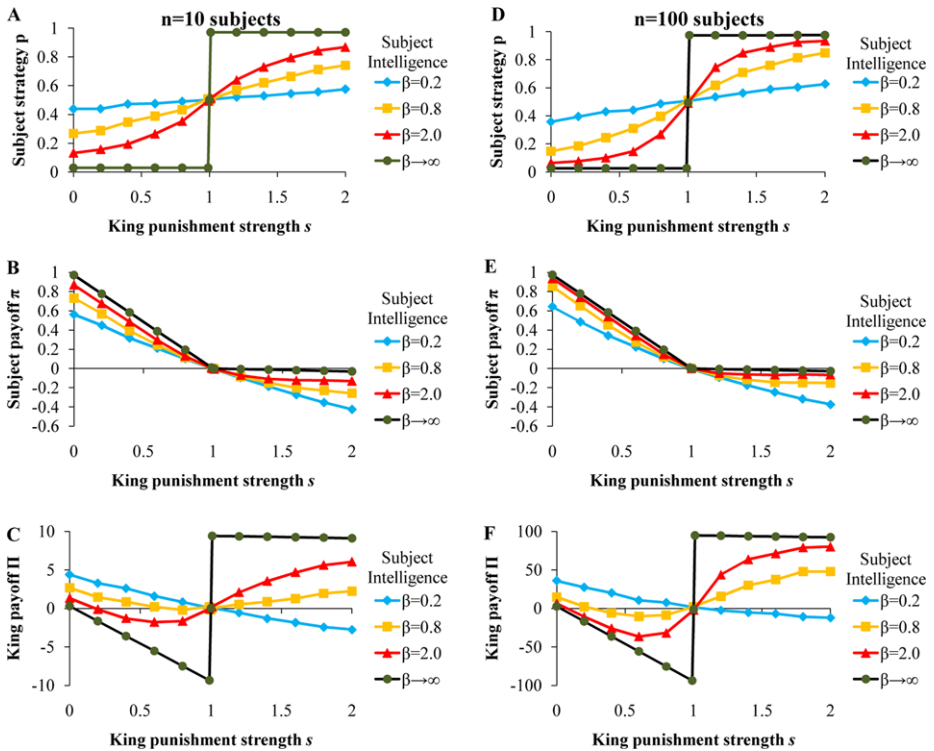


Fig. 1 Subject tribute payment p (Row 1), subject payoff π (Row 2), and king payoff Π (Row 3) in a population of $n = 10$ (Column 1) or $n = 100$ (Column 2) subjects evolving under a fixed level of king punishment s

amount. For low values of punishment strength s , king payoff decreases as punishment strength increases. Here most subjects pay little tribute, and so a large amount of punishment is carried out—therefore the increased cost of punishment when s increases outweighs the gains from the resulting small increase in tribute payment p . Eventually, the balance between this cost and benefit of increasing s shifts, and at higher s values the relationship between king payoff and punishment strength flips. After some threshold value of s , king payoff is increasing in punishment strength (for $\beta \rightarrow \infty$, this transition happens precisely at $s = 1$). Thus, as long as subjects are sufficiently intelligent, kings do best by punishing severely.

Finally, we see that population size has little effect. For both $n = 10$ and $n = 100$, the results are very similar. Thus we will use $n = 10$ for computational tractability when considering the case of multiple kingdoms.

3.3 Evolutionary Dynamics: Multiple Kingdoms

We now allow kings to learn as well as subjects, and turn our attention to the coevolution of subjects and kings across multiple kingdoms. We consider $N = 10$ kingdoms, each consisting of $n = 10$ subjects and one king. We examine the average king’s strategy s in addition to the average subject’s strategy p , subject’s payoff π , and king’s payoff Π . We ask how each quantity varies as a function of subject intelligence $\beta \in \{0.2, 0.8, 2, \infty\}$ and the probability of a king update round $\theta \in [0.001, 1)$. In each simulation, we fix subject mutation

Table 1 Summary of coevolutionary outcomes

		Subject intelligence β	
		Low	High
Subject update probability $(1 - \theta)$	Low	Intermediate p , small s	Small p , small s
	High	Intermediate p , small s	Large p , large s

rate $\mu_s = 0.05$ and king mutation rate $\mu_k = 0.05$, and initialize each subject with a randomly drawn strategy from the interval $[0, 1]$ and each king with a randomly drawn strategy from the interval $[0, 3]$. We then simulate for 10^7 strategy updates and calculate the values of s, p, π , and Π averaged over the second half of the simulation. We repeat this same analysis for both probabilistic king learning ($\beta_K = 0.1$) and deterministic king learning ($\beta_K \rightarrow \infty$). Recall that subjects learn only from other subjects in the same kingdom, while kings learn from the kings of other kingdoms.

The results are shown in Fig. 3 and summarized in Table 1. For both probabilistic and deterministic king learning, we see that natural selection favors strong king punishment and high subject tribute payments when (i) subject intelligence is sufficiently high and (ii) subjects update sufficiently rapidly relative to leaders (i.e., leader update probability θ is sufficiently small). As we saw in the single kingdom case above, subject intelligence is needed in order for the punishment incentive to effectively remodel subject behavior.

The additional condition that subjects must update more quickly than kings occurs for the following reason: Imagine a king who punishes little and as a result rules over subjects who pay little tribute. Now imagine that the king experiments with a stronger punishment strategy. At first, he earns less than he was earning before—he incurs greater punishment costs when punishing his subjects, most of whom are noncompliant. It is only over time, as the selection gradient on subjects set up by the stronger punishment leads to higher subject tribute payments, that the king reaps a benefit from his increased punishment. Thus, if kings update too quickly (i.e., subjects update too slowly), the kings will abandon stronger punishment before it has a chance to have an effect. To be effective, kings must be resolute. By the same token, the smarter the subjects are, the more quickly they respond to punishment; and therefore higher subject intelligence allows for the success of punishment under faster king updating.

Considering king payoffs, we see that kings benefit from having subjects who update accurately and frequently. When subjects are intelligent but slow, kings do very poorly. As explained above, the kings cannot effectively learn to punish when subjects update slowly; and because the subjects are intelligent, they therefore learn to pay little tribute.

Considering subject payoffs, we again see an interaction between the effects of intelligence and update speed. When subjects update quickly compared to kings, the subjects are better off with low intelligence. That is because if subject intelligence is high in this regime, kings can effectively impose punishments, forcing the subjects to pay tribute. But if the subjects are sufficiently confused, then the kings abandon punishment, and the subjects wind up paying less tribute. When subjects update slowly relative to the kings, the opposite is true: subject payoffs are increasing in intelligence. As discussed above, this is because kings do not learn to impose punishments in this regime, and so intelligent subjects learn to pay less tribute.

Next, we demonstrate that these results are robust to varying our (arbitrarily imposed) upper limit on king punishment strength s . In Fig. 2, we use the interval $s \in [0, 3]$. We now replicate the same simulations using $s \in [0, 10]$ and get very similar results (see Fig. 3).

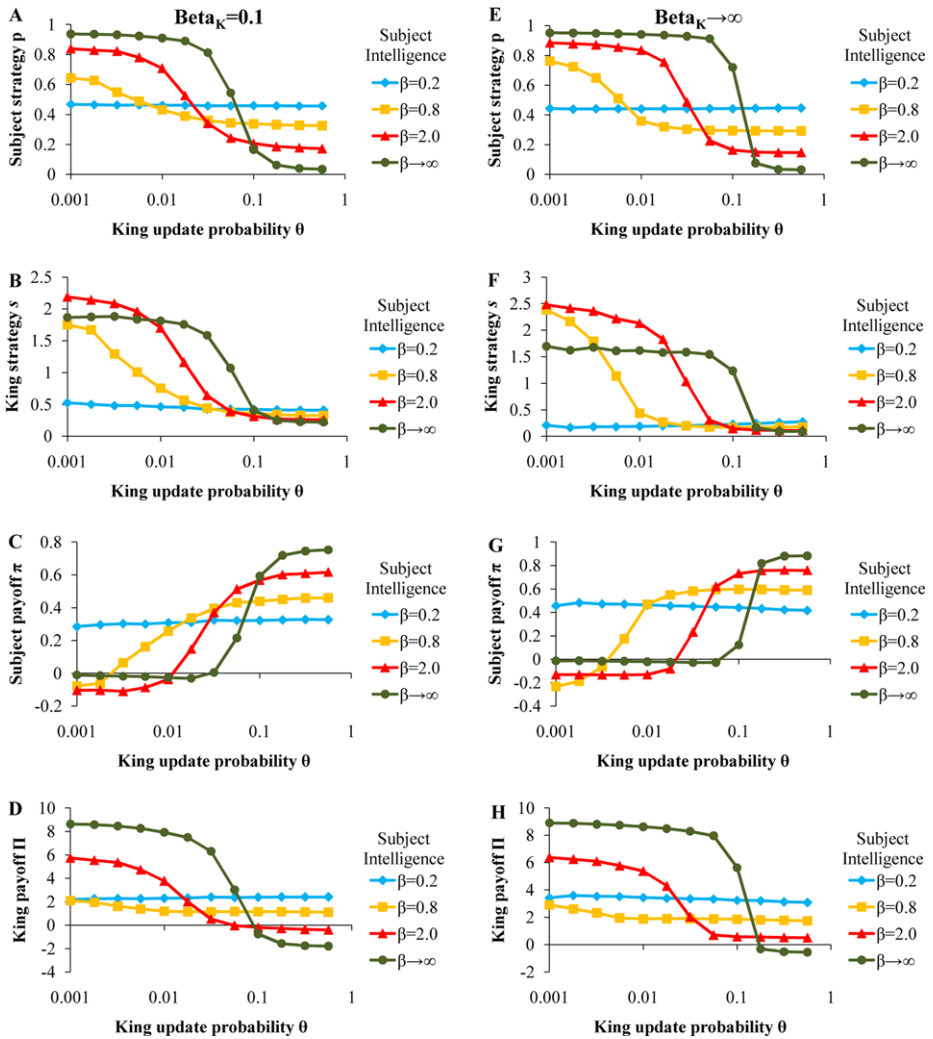


Fig. 2 King punishment strength s (Row 1), subject tribute payment p (Row 2), subject payoff π (Row 3), and king payoff Π (Row 4) resulting from the coevolution of kings and subjects in a world of 10 kingdoms, each made up of $n = 10$ subjects, with kings using probabilistic learning (Column 1) or deterministic learning (Column 2)

Finally, we explore the consequences of inter-kingdom learning among subjects. Thus far, we have only allowed subjects to learn from other subjects within the same kingdom. We now ask what happens if instead a given subject compares her payoff to a randomly chosen other subject from the entire population with probability ϕ , or from within the same kingdom with probability $1 - \phi$. Thus, ϕ is a measure of the “global awareness” of subjects. When $\phi = 0$, subjects are completely unaware of the world outside their own kingdom. As ϕ increases, subjects become more and more informed about the behavior of subjects in other kingdoms. To investigate the effect of increasing ϕ , we choose two sets of parameters from Fig. 2 which were favorable for coercion ($n = N = 10$, $\theta = 0.01$, $\mu_s = \mu_k = 0.05$, with either probabilistic learning $\beta = 2$, $\beta_K = 1$ or deterministic learning, $\beta \rightarrow \infty$, $\beta_K \rightarrow \infty$) and

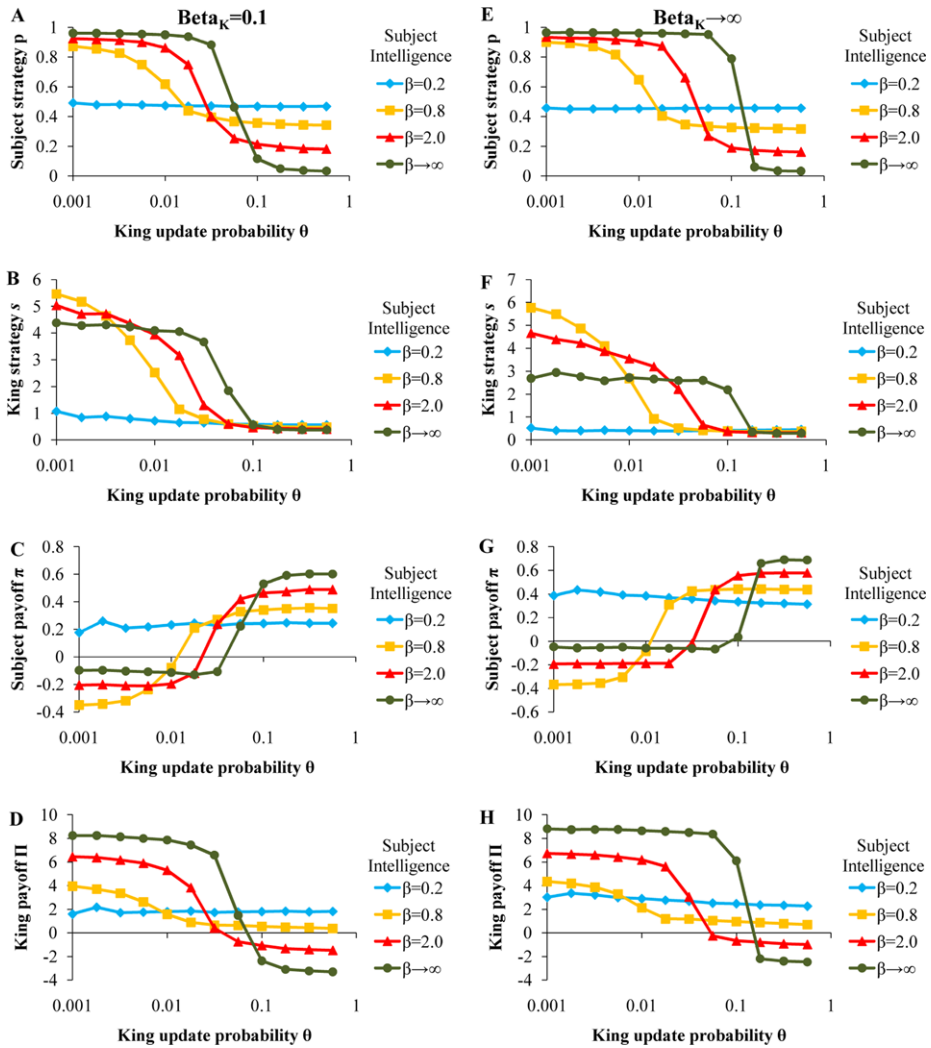
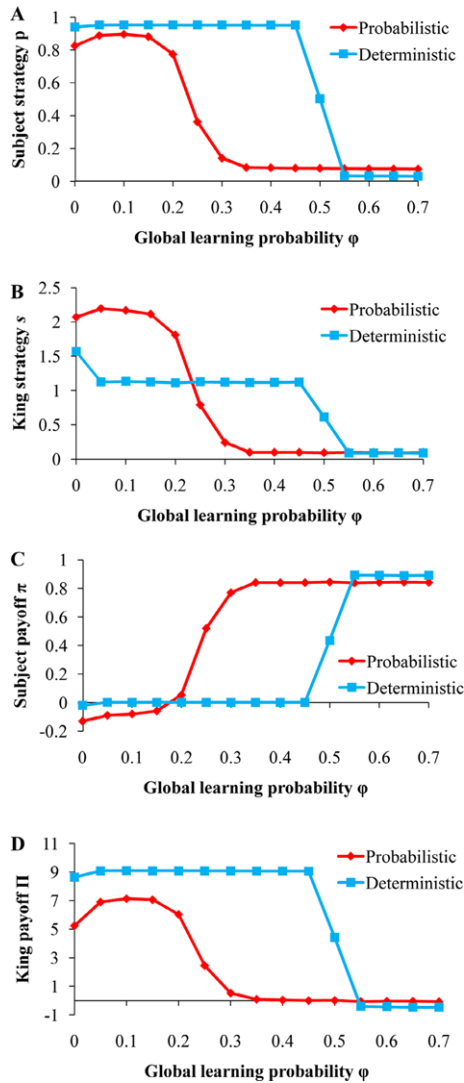


Fig. 3 Demonstration of the robustness to changing the maximum punishment strength s . This figure replicates the results from Fig. 2 using a maximum of $s = 10$ rather than the value of $s = 3$ used previously

simulate for $\phi \in [0, 0.7]$. The results are shown in Fig. 4. We see that inter-kingdom learning has little effect until a critical threshold is reached, at which point punishment collapses and subjects stop paying tribute. When inter-kingdom learning is sufficiently frequent, a social dilemma arises amongst the kings. If some kings pay to punish while others do not, the subjects of the punishing kings will see the subjects of the lenient kings earning high payoffs, and stop paying tribute (even though as a result they get punished); and as a result, the kings will evolve to stop punishing. We also observe that higher fidelity learning allows coercion to be maintained in the face of a higher level of global awareness.

Fig. 4 Allowing subjects to learn from subjects in other kingdoms undermines coercion. Shown is king punishment strength s (A), subject tribute payment p (B), subject payoff π (C), and king payoff Π (D) resulting from the coevolution of kings and subjects for various probabilities ϕ of global (inter-kingdom) learning among subjects. We consider probabilistic updating (*diamonds*; $\beta = 2$, $\beta_K = 1$) and deterministic updating (*squares*; $\beta \rightarrow \infty$, $\beta_K \rightarrow \infty$), both using king update probability $\theta = 0.01$



4 Discussion

Here we have demonstrated how natural selection can lead to the emergence of institutional punishment in connection with coercion rather than cooperation. Tyrannical rulers who use punishment to extort tribute from their subjects may outperform more benevolent leaders. As a result, harsh state-inflicted punishment may spread across societies.

This process of inter-state learning bears some similarities to models of group (or “multilevel”) selection [3, 15, 16, 22, 23, 29, 33, 37–39]. Our model involves a co-evolutionary process with selection in some sense occurring at two “levels”: subjects learn from subjects, and kings learn from kings. The framework we introduce, however, has important differences from classical group selection models. Our model does not require massive intergroup conflicts in which one society is entirely replaced by another. Instead, all that is

required is one individual ruler deciding to imitate another. A single leader or governing group changing policy then leads to most individuals in the group changing their strategy as a result of new top-down incentivizes, rather than through mass extermination (or imitation) as in models based on inter-group conflict. Thus large-scale societal change can occur based on the decision of only a small number of individuals in positions of power. Furthermore, while retaliation can undermine the effectiveness of punishment in a group selection framework [14], retaliation is much less of a concern in the asymmetric interactions between state and individual.

In our model, leaders are selfish individuals seeking only to maximize their own payoffs at the expense of their subjects. This is, of course, an extreme position. A state's income is often used to provision public goods, providing a benefit to the citizens of the state. We show that even in the absence of such benefits, punishment can prove effective. Furthermore, we show that in our framework, even selfish leaders have an incentive to educate their citizens. An intelligent, well-informed, and frequently updating citizenry responds more quickly and fully to coercive incentives imposed by their leaders. The necessity for slowly updating kings may also help explain why successful institutions are slow to change their policies and why most states are so bureaucratic. It is necessary for a state to be steadfast in its implementation of punishment in order to succeed in coercing its subjects.

One feature of society which our model leaves out is the possibility of rebellion and revolution. Educating the public may lead to discontent and ultimately the downfall of a despotic state. In line with this reasoning, we find that "global awareness" among the subjects undermines kings' ability to coerce. Thus although kings do best when their subjects are intelligent in the sense of being accurate at assessing payoffs, this is only true if subjects' awareness is largely limited to others within the same kingdom. In our model, coercion relies on censorship and "closed minded" subjects who either ignore or are uninformed about norms and behavior outside of their own nation. This result is suggestive of recent political events in which the internet and social media, which produce greater global awareness, have helped precipitate the overthrow of oppressive governments.

Related to this issue is the potential effect of king "debt ceilings." In the current model, kings can always punish, regardless of their income. If kings were unable to punish once their earnings drop below some threshold, this could create the potential for collective action by subjects: if enough subjects refused to pay tribute (and thereby expose themselves to "potential punishment"), the king could be "starved" out of power. Expanding the model to include such phenomena is a fruitful direction for future research, as is considering more democratic institutions where citizens can vote tyrannical leaders out of power.

Viewing the issue of intelligence and update speed from the perspective of the subjects provides another interesting result. As we saw in the one kingdom case, for any given level of punishment, a subject would increase her personal payoff by being more intelligent. But once kings are also able to learn, the coevolutionary dynamic leads to better subject outcomes when subjects are confused. Thus there is a tension between individual and collective interests, and confusion is a form of public good for the subjects.

Another promising avenue for future research involves behavioral experiments [24] investigating institutional punishment. Important previous experiments involving endogenous choice have explored the emergence of symmetric institutions, with public goods game players choosing to play in games with or without sanctions [9, 28, 32]. Another paper explored a symmetric public goods game with asymmetric punishment—only one group member was allowed to punish [19]. Under those circumstances, punishment was largely used to promote cooperation. Our model suggests an alternate paradigm where the party that can punish also benefits disproportionately from the contributions of others. The only partially asymmetric

setup of [19] may prime subjects to see the game as cooperative. We would predict that in the entirely asymmetric scenario modeled here, human subjects would be more likely to use punishment coercively when put in the role of king. Exploring this and other related questions experimentally will help shed light on the origins of coercive institutional punishment.

In summary, we have shown one simple method by which institutional punishment could have evolved through inter-state learning. We hope that this work will inspire further exploration of institutional punishment from an evolutionary perspective, both theoretically and experimentally. Doing so is essential for understanding the world in which we live, which is dominated by strong institutions which punish law-breakers.

Acknowledgements We thank Martin Nowak and the Program for Evolutionary Dynamics at Harvard University for useful discussions which inspired this work, and Feng Fu, Benedikt Herrmann, and Michael Manapat, as well as the editor and three anonymous referees, for helpful comments on the manuscript. D.G.R. is supported by a grant from the John Templeton Foundation.

References

1. Abbink K, Herrmann B (2009) The moral costs of nastiness. Discussion Paper, Cent Decis Res Exp Econ, University of Nottingham
2. Boyd R, Richerson PJ (1992) Punishment allows the evolution of cooperation (or anything else) in sizeable groups. *Ethol Sociobiol* 13:171–195
3. Boyd R, Gintis H, Bowles S, Richerson PJ (2003) The evolution of altruistic punishment. *Proc Natl Acad Sci USA* 100:3531–3535
4. Burnham T, Johnson DDP (2005) The evolutionary and biological logic of human cooperation. *Anal Krit* 27:113–135
5. Clutton-Brock TH, Parker GA (1995) Punishment in animal societies. *Nature* 373:209–216
6. Dawes CT, Fowler JH, Johnson T, McElreath R, Smirnov O (2007) *Nat Lett* 446:794–796
7. Dreber A, Rand DG, Fudenberg D, Nowak MA (2008) Winners don't punish. *Nat Lett* 452:348–351
8. Fehr E, Gächter S (2002) Altruistic punishment in humans. *Nature* 415:137–140
9. Gurerk O, Irlenbusch B, Rockenbach B (2006) The competitive advantage of sanctioning institutions. *Science* 312:108–111
10. Hauert C, Traulsen A, Brandt H, Nowak MA, Sigmund K (2007) Via freedom to coercion: the emergence of costly punishment. *Science* 316:1905–1907
11. Helbing D, Szolnoki A, Perc M, Szabo G (2010) Evolutionary establishment of moral and double moral standards through spatial interactions. *PLoS Comput Biol*. doi:[10.1371/journal.pcbi.1000758](https://doi.org/10.1371/journal.pcbi.1000758)
12. Henrich J, Boyd R (2001) Why people punish defectors: weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. *J Theor Biol* 208:79–89
13. Herrmann B, Thoni C, Gächter S (2008) Antisocial punishment across societies. *Science* 319:1362–1367
14. Janssen MA, Bushman C (2008) Evolution of cooperation and altruistic punishment when retaliation is possible. *J Theor Biol*. doi:[10.1016/j.jtbi.2008.06.017](https://doi.org/10.1016/j.jtbi.2008.06.017)
15. Keller L (1999) *Levels of selection in evolution*. Princeton Univ Press, Princeton
16. Michod RE (1999) *Darwinian dynamics*. Princeton Univ Press, Princeton
17. Nakamaru M, Iwasa Y (2005) The evolution of altruism by costly punishment in lattice-structured populations: score-dependent viability versus score-dependent fertility. *Evol Ecol Res* 7:853–870
18. Nakamaru M, Iwasa Y (2006) The coevolution of altruism and punishment: role of the selfish punisher. *J Theor Biol* 2450:475–488
19. O'Gorman R, Henrich J, Van Vugt M (2009) Constraining free riding in public goods games: designated solitary punishers can sustain human cooperation. *Proc R Soc B* 276:323–329
20. Ostrom E, Walker J, Gardner R (1992) Covenants with and without a sword: self-governance is possible. *Am Polit Sci Rev* 86:404–417
21. Panchanathan K, Boyd R (2004) Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature* 432:499–502
22. Paulsson J (2002) Multileveled selection on plasmid replication. *J Genet* 161:1373–1384
23. Rainey PB, Rainey K (2003) Evolution of cooperation and conflict in experimental bacterial populations. *Nature* 425:72–74
24. Rand DG (2011) The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. *J Theor Biol*. doi:[10.1016/j.jtbi.2011.03.004](https://doi.org/10.1016/j.jtbi.2011.03.004)

25. Rand DG, Dreber A, Ellingsen T, Fudenberg D, Nowak MA (2009) Positive interactions promote public cooperation. *Science* 325:1272–1275
26. Rand DG, Ohtsuki H, Nowak MA (2009) Direct reciprocity with costly punishment: generous tit-for-tat prevails. *J Theor Biol* 256:45–47
27. Rand DG, Armao JJ, Nakamaru M, Ohtsuki H (2010) Anti-social punishment can prevent the co-evolution of punishment and cooperation. *J Theor Biol* 265:624–632
28. Rockenbach B, Milinski M (2006) The efficient interaction of indirect reciprocity and costly punishment. *Nature* 444:718–723
29. Rogers AR (1990) Group selection by selective emigration: the effects of migration and kin structure. *Am Nat* 135:398–413
30. Sigmund K (2007) Punish or perish? Retaliation and collaboration among humans. *Trends Ecol Evol* 22:593–600
31. Sigmund K, De Silva H, Traulsen A, Christoph H (2010) Social learning promotes institutions for governing the commons. *Nature* 466:861–863
32. Sutter M, Haigner S, Kocher MG (2010) Choosing the carrot or the stick? Endogenous institutional choice in social dilemma situations. *Rev Econ Stud* 77:1540–1566
33. Taylor PD, Wilson DS (1988) A mathematical model for altruism in haystacks. *Evolution* 42:193–197
34. Traulsen A, Pacheco JM, Nowak MA (2007) Pairwise comparison and selection temperature in evolutionary game dynamics. *J Theor Biol* 246:522–529
35. Traulsen A, Hauert C, de Silva H, Nowak MA, Sigmund K (2009) Exploration dynamics in evolutionary games. *Proc Natl Acad Sci USA* 106:709–712
36. Wang J, Wu B, Ho DWC, Wang L (2011) Evolution of cooperation in multilevel public goods games with community structures. *Europhys Lett* 93:58001
37. Williams GC, Williams DC (1957) Natural selection of individually harmful social adaptations among sibs with special reference social insects. *Evolution* 11:32–39
38. Wilson EO, Holldobler B (2005) Eusociality: origin and consequences. *Proc Natl Acad Sci USA* 102:7411–7414
39. Wilson DS, Sober E (1994) Reintroducing group selection to the human behavioral sciences. *Behav Brain Sci* 17:585–608
40. Yamagishi T (1986) The provision of a sanctioning system as a public good. *J Pers Soc Psychol* 51:110–116