Episteme, 10, 4 (2013) 403–416 © Cambridge University Press
doi:10.1017/epi.2013.34

POLICING EPISTEMIC COMMUNITIES

JUSTIN P. BRUNER
brunerj@uci.edu

ABSTRACT

I examine how particular social arrangements and incentive structures encourage the honest reporting of experimental results and minimize fraudulent scientific work. In particular I investigate how epistemic communities can achieve this goal by promoting members to police the community. Using some basic tools from game theory, I explore a simple model in which scientists both conduct research and have the option of investigating the findings of their peers. I find that this system of peer policing can in many cases ensure high levels of honesty.

1. INTRODUCTION

Science is rightfully lauded as one of the most successful human enterprises. Yet what accounts for its immense success? Recently, philosophers of science have attempted to explain the vast success of science with the help of economic models. Kitcher (1990) and Strevens (2001), for example, both rely on simple models to justify various norms in science, such as the priority rule. A common assumption in their work (as well as other's work in the economics of science literature) is that individual scientists are self-interested. That is, individual scientists are not modeled as egoless truth seekers, but reputation maximizers. Scientists are by and large motivated by their reputation or standing in a given academic community, and this in part determines which research project they devote time to, who they cite, etc. By each scientist pursuing her self-interests, the community's projects are advanced.

Yet conceiving of scientists as self-interested reputation maximizers comes with its own set of problems not addressed in these models. If scientists can boost their reputation by publishing novel and exciting findings, then scientists face the temptation of fabricating data. In fact, since a scientist's reputation is so closely tied to her publication record, scientists can be easily tempted to engage in some form of misbehavior (outright fraud, plagiarism, self-plagiarism). Unfortunately, reviews of the medical and biological sciences have revealed that the proportion of articles retracted due to fraud has "increased 10-fold since 1975" (Frang et al. 2012). This suggests that while the incentive structure scientists face can often promote competition between scientists that results in scientific breakthroughs and discoveries, it also promotes fraud and careless, error-prone work.

Most of the hitherto work in the social epistemology of science has overlooked these perverse incentives in science. A natural question is then how do scientific communities minimize the amount of fraud and misbehavior? I explore one such means of ensuring minimal levels of fraud and corner-cutting in science. Specifically, I investigate the possibility of peer-policing in science, that is, arrangements in which a scientist's peers are

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

responsible for identifying errors in her work in addition to ensuring she does not get away with fraud. This of course neglects other mechanisms at our disposal, such as the peer review process, in which journal editors and referees jointly assess papers. This latter arrangement has been briefly studied in Cox and Goldman (1994), although not with an eye toward identifying fraud.[1] I instead investigate the conditions under which a decentralized community of scientists can successfully police themselves.

In the course of this paper I will rely on some basic tools from game theory and evolutionary game theory. This allows us to precisely capture the strategic situation scientists find themselves in and study the behavior of the community over time. We utilize a low-rationality approach, where it is assumed that individuals are both myopic and have a propensity to imitate the behavior of those more successful than themselves. This allows us to carve out a class of incentive schemes that promote peer-policing. Essential to all of this will be the use of punishment in science. In the next section we make explicit connections to the evolution of altruism literature, where punishment is often invoked to sustain prosocial behavior. We then construct our game theoretic model. We find in section 3 that if agents are rewarded for punishing their peers then science can truly be thought of as a self-policing epistemic community – frauds and cheats will be kept at bay due to the vigilance of their peers. Section 4 concludes with a brief discussion of how to interpret my results, discusses the normative implications of my model and gestures toward future work.

2. PUNISHMENT IN BOTH SCIENCE AND NATURE

We begin with a brief discussion of the evolution of altruism. In the course of this discussion we uncover some mechanisms that we will ultimately utilize in our analysis of the social structure of science. The connection between these two disparate topics is that in both contexts punishment can be administered by the community at large in order to ensure that individuals behave appropriately. In the case of altruism, community enforcement helps ensure that members behave cooperatively. In science, punishment can ensure that scientists adhere to the standards and norms of the discipline. Thus punishment plays a vital role in both human cultural groups and scientific communities. However, while punishment has been much explored in the case of cooperation, no such similar investigation has taken place in the case of science. In the course of this section we further discuss the use of punishment in the evolution of cooperation and further characterize the role punishment plays in science.

Altruism is often taken to be an evolutionary puzzle. Altruistic behavior is typically defined as behavior which decreases the fitness of the altruist while simultaneously increasing the fitness of the recipient of the altruistic behavior. Evolution would thus be expected to select against altruistic behavior. Yet the fact of the matter is that altruistic behavior is commonly observed, in both human and non-human organisms. This large gap between the facts and the theory has generated a cottage industry intent on providing a satisfactory theoretical explanation for these high levels of altruistic behavior. What has been uncovered in the last 40 to 50 years is that there are a number of different mechanisms that can sustain altruistic behavior.

1 It should also be noted that Cox and Goldman explore a different sort of epistemic community, discussing the role newspaper editors play in journalism.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

One such mechanism is that of *altruistic punishment*.[2] The idea is a deceptively simple one. We introduce "punishers" to the population who cooperate with others and punish those who fail to behave altruistically. In other words, these punishers inflict a cost upon defectors who fail to cooperate with their peers. If this cost is high enough, then it now becomes individually rational for members of the community to behave altruistically, so as to avoid the costs associated with defecting. Thus, when there are enough punishers in the population, altruistic behavior can be reliably sustained.

This brand of punishment is often referred to as altruistic punishment for two reasons. First, punishment of this sort is altruistic in the sense that it is done in an attempt to enforce and stabilize pro-social norms in the population (although it should be noted that this sort of punishment can likewise be utilized to enforce much more nefarious norms as well). Second, it is altruistic in the sense that the act of punishing another is often costly. Mathematical models that incorporate punishment often assume that there is a cost, albeit a small one, associated with punishing another. This is reasonable if one assumes that there is a small chance of retaliation, or that it is cognitively taxing to keep track of who is a cooperator or defector.

That altruistic punishment is costly to the individual means it is never individually rational to punish. Ergo while cooperation can be sustained through punishment, the question of how to sustain high levels of punishment must be addressed.[3] Resolving this second-order free-rider problem has garnered much attention in recent years. Many proposals claim that by incorporating other social mechanisms such as group selection (Boyd et al. 2003) or image scoring (Panchanathan and Boyd 2004) along with punishment it is possible to sustain both cooperation and punishment.

Recently both biologists and philosophers of biology have considered a different route. Instead, they argue, there is no puzzle as to why punishment evolved because the second-order free-rider problem is predicated on a fallacious premise – that punishing another is individually costly. If punishment somehow directly benefits the punisher it is no great mystery that punishment could evolve. Since this brand of punishment is not costly and thus not altruistic, we'll refer to it as *opportunistic punishment*.[4] Examples of opportunistic punishment in the biological sciences abound. For instance, Rohwer (2006) notes that in chimpanzees the alpha male has incentive to punish other group members when this punishment reinforces the status quo, a state of affairs the alpha male strongly prefers. Cant and Johnstone (2006) note that opportunistic punishment occurs among fish when selecting which smaller "cleaner fish" they will associate with. Additionally, Nakao and Machery (2012) provide a detailed account of the important role opportunistic punishment has had in a number of other organisms.

It is not a new realization that the problems associated with punishment can be avoided if punishment is opportunistic. David Hull understood the potential of opportunistic punishment a decade before it was noted in the biological literature. Interestingly, he was not

2 See Sripada (2005) for a persuasive argument that altruistic punishment, and not reciprocity, plays a pivotal role in the evolution of cooperation as well as the enforcement of social norms. See also Boyd and Richerson (1992).

3 Andrew Colman puts it particularly elegantly: "we seem to have replaced the problem of explaining cooperation with that of explaining altruistic punishment" (Colman 2006).

4 This is my terminology. In the literature opportunistic punishment often goes by the name of self-serving punishment.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

concerned with altruism, but how to ensure high levels of quality in science and minimize the occurrence of erroneous results. For example, he notes how the incentive structure in science is such that individuals often get credit for pointing out the failures of their peer's papers.

> Scientists get very little credit for replicating other scientists' experiments, but they do get credit for discovering mistakes in the work of others, especially if this research is taking place in one of the "hot" areas of science. The rush to publish, when properly constrained, increases the pace of science. The monitor on this pace is the punishment meted out to those scientists who produce unreliable work. Thus, science can be viewed as a self-policing system of mutual exploitation – or cooperation if one prefers. (Hull 1996a)

Scientists have incentive to cut corners with their research, or even worse – fabricate data. Hull suggests that one natural means of ensuring that this behavior does not become too rampant is that the incentive structure scientists face encourages them to point out the errors and mistakes in past work, or even uncover fraud!

The remainder of this paper will further explore this idea in an attempt to determine how feasible and how successful such an arrangement can be. In particular, I construct a model of a decentralized community of boundedly rational scientists to demonstrate how effective opportunistic punishment can be.

3. THE MODEL AND RESULTS

We consider a population of scientists all engaged in some form of scientific inquiry. In each time period, scientists work on a project. We assume for simplicity that no two scientists can work on the same project. At the start of the time period, the scientist is ignorant of how successful this project will be. In other words, she does not know if completing the project will lead to interesting and novel results, or will lead to rather trivial results. The only way to uncover the "worth" of the project is to actually engage in research. By the end of the time period, the scientist has worked on the project and successfully uncovered the true value of the project.

We can represent the importance or value of a completed project as a number. The higher this number is, the more valuable the completed project is. In our model, the worth of a scientist's project is determined by a random draw from a pre-determined probability density function. We work with the idealizing assumption that all scientists pull from the same function (in other words, all scientists are equally talented). Additionally, scientists have the option of truthfully or fallaciously reporting the 'worth' of their project. Since only the scientist conducting the research initially has access to the number given by the distribution, her peers cannot observe if she has chosen to accurately report this value to the population or "embellish" and add a small term, θ, to the importance of her work.

We restrict our study to three simple strategies: honestly report (H), police (P) and cheat/cut-corners (C). The strategic form of the game is represented in Table 1. We refer to this game as the Merit-Demerit game. Honestly report is a "safe strategy" in the sense that it always guarantees an expected payoff of μ (the mean of the distribution). The two remaining strategies are a bit riskier, because how well they perform is contingent

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

Table 1. The Merit-Demerit Game. μ is the average draw from the distribution, c is the cost associated with policing, b is the benefit to the punisher of punishing a cheater, θ is the benefit of cheating and p is the punishment associated with cheating and getting caught. We assume that all of these values are greater than zero.

	H	*P*	*C*
H	μ, μ	μ, μ − c	μ, μ + θ
P	μ − c, μ	μ − c, μ − c	μ − c + b, μ + θ− p
C	μ + θ, μ	μ + θ− p, μ − c + b	μ + θ, μ +θ

on what their counterpart in the game does. A cheater/corner-cutter would obviously prefer to interact with a fellow cheater/corner-cutter or honest type than a policer. Likewise, since a policer only gets acknowledgement for catching someone, she strongly prefers to interact with a cheater/corner-cutter.

When $\theta < p$ and $\theta + c < b$ there exist no pure-strategy equilibria in the Merit-Demerit game. These parameter restrictions are reasonable. $\theta < p$ simply states that the punishment associated with getting caught is greater than the benefit of cheating, while $\theta + c < b$ ensures that the cost of patrolling the community is at least θ less than the benefit of catching someone. The mixed-strategy equilibrium of this game is $(1 - \frac{\theta}{p} - \frac{c}{b}, \frac{\theta}{p}, \frac{c}{b})$.[5] From the mixed-strategy equilibrium alone we can already gain some insight as to what sort of payoffs should be in place to ensure mostly honest behavior. Note that as p, the punishment meted out to cheaters, and b, the benefit of catching a cheat increase, the proportion of honest types increases.

While we have identified a mixed-strategy Nash equilibrium, we still need assurance that boundedly rational agents will converge to this equilibrium. To assess whether this is the case we turn to evolutionary game theory. We rely for the most part on the replicator dynamics (Taylor and Jonker 1979). This is a standard mathematical model of biological evolution. Surprisingly, the replicator dynamics are formally equivalent to many models of social evolution as well. For example, many forms of cultural evolution involving differential imitation (i.e., situations in which one imitates the behavior of more successful peers) turn out to be mathematically equivalent to the replicator dynamics.[6] This in part justifies our use of a model of biological evolution in the study of human behavior. Since there are strong connections between the replicator dynamics and standard models of cultural evolution, exploring the behavior of the population under the replicator dynamics provides us with some insight as to how these cultural groups would actually change over time. Additionally, in many instances the replicator dynamics and the best response dynamics yield the same qualitative results.[7] The logic of the replicator dynamics is simple enough. Strategies that on average do better than the population as a whole

5 This of course holds when $\frac{\theta}{p} + \frac{c}{b} < 1$ (and neither p nor b are equal to zero).

6 For more on the relationship between the replicator dynamics and cultural evolution see Sandholm (2010). See also Schlag (1997).

7 In the best response dynamic at each time period a proportion of the population is selected to "best respond" to the current population. It should be pointed out that it is well known that best response dynamics and the replicator dynamics diverge, especially in the case of "spiraling." Thus our results under the replicator dynamics (especially those in Figures 1a and 2a) may not hold under the best response dynamics.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

increase in frequency, while strategies that on average do worse than the population as a whole decrease in frequency. The continuous-time replicator dynamics is

$$\dot{x}_k = x_k(\overline{F_k} - \overline{F})$$

where x_k is the proportion of individuals playing strategy k, $\overline{F_k}$ is the average fitness of strategy k and $\overline{F}$ is the average fitness in the population.[8] The replicator dynamics allow us to track the population over time from a randomly selected starting point.

One may wonder why the behavior of otherwise rational and intelligent scientists would be explored with a low-rationality model common to evolutionary game theory. After all, while the replicator dynamics may be an apt model of non-human animals, surely humans – and especially scientists! – deserve a more sophisticated modeling treatment. While scientists are sophisticated learners in their laboratories, there are a number of good reasons why we should take the low-rationality approach when modeling their behavior. First I am interested in identifying the minimal levels of rationality required of our agents in order to ensure that the community as a whole remains honest. If we can demonstrate that simple learners can successfully police the community, then our work is done. Second, many anthropologists and social scientists remain convinced that rather simple rules are utilized by humans when determining how to update their behavior. The conformist bias is one such rule (Richerson and Boyd 2005), and instructs individuals to adopt behavior that is prevalent in the population (as opposed to behavior that is well-adapted). How sophisticated human behavior really is of course remains an empirical question, but we nonetheless have some good reason to believe a low-rationality model may accurately capture human behavior. Finally, the low-rationality model of behavior utilized in this paper allows us to side-step many of the traditional criticisms often levied at the economic approach to human behavior. Complete information, perfect recall and knowledge of equilibrium behavior are not necessary assumptions.[9]

We find that our epistemic community can successfully self-police, but only when policing is rewarded, i.e., when punishment is opportunistic. Figure 1 compares two cases. In Fig. 1(a), we examine the Merit-Demerit game when there is both a benefit associated with catching a cheat ($b > 0$) and a cost associated with being caught ($p > 0$). In this case punishment is not necessarily costly to the punisher. Whether punishment is altruistic or opportunistic in the Merit-Demerit game hinges on the composition of the population. More specifically, policing is individually rational if the proportion of cheaters in the population is sufficiently high. If there are too few cheaters in the population, policers exert effort patrolling their ranks for naught – they will spend precious resources investigating the work of others and rarely uncover misbehavior.

Thus our epistemic community is characterized by the cyclic behavior seen in Fig. 1 (a). When all are honest, cheaters abound. As the proportion of cheaters increases, policing becomes increasingly appealing, ushering in policers. Once the number of cheaters dips below some threshold value it no longer pays to police, and honest types increase in

8 $\overline{F} = \Sigma_{j=1}^{n} x_j \overline{F_j}$ and $\overline{F_k} = \Sigma_{j=1}^{n} x_j F_{kj}$, where F_{kj}, is the payoff strategy k receives upon interacting with strategy j.

9 Of course exploring issues of fraud and policing with help from a more "traditional" economic model is a worthwhile endeavor as well. Wible (1992) studies both fraud and replication in science using a model in which economic agents determine their course of action by solving an optimization problem.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

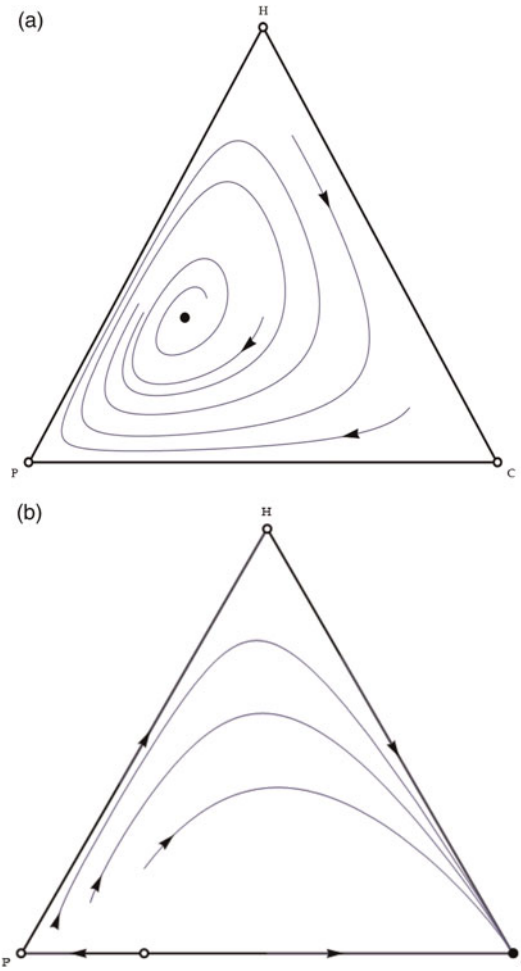


Fig. 1. (colour online) (a) Dynamics of the Merit-Demerit game with $\mu = 10$, $\theta = 1$, $p = 2$, $b = 3$ and $c = 0.5$. The polymorphic equilibrium (1/3, 1/2, 1/6) attracts all interior points of the simplex. (b) Dynamics of the Merit-Demerit game with altruistic punishment. With parameter values $\mu = 10$, $\theta = 1$, $p = 2$, $b = 0$ and $c = 0.5$, all interior points lead to cheat.

frequency, which in turn opens the door for cheaters to re-emerge. This cat and mouse game continues as the population spirals to the equilibrium in which the majority of agents either police or honestly report their results. At the polymorphic equilibrium in Fig. 1(a) half of the agents police, one-third honestly report and only one-sixth cheat.[10]

By adjusting the values in the payoff table we can increase the proportion of honest types at the polymorphic equilibrium. It is desirable to have low levels of both cheaters and policers since policers waste time and energy that could have been spent conducting research. For example, if we change p and b to $p = 4$, $b = 5$ while holding everything else

10 In this case the polymorphic equilibrium is (almost) a global attractor. All points in the interior of the simplex are drawn to the polymorphic equilibrium.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

constant, the population is drawn to a polymorphic equilibrium where 25% police, 65% are honest types, and 10% cheat. Furthermore, bumping p to 6 and b to 7 yields even more desirable results. In this case 16.7% police, 76.2% are honest, and 7.1% cheat.[11]

Punishment fails to be effective if it is too costly. In Fig. 1(b), we examine the Merit-Demerit game with altruistic punishment.[12] Recall that with this type of punishment it is never the case that one directly benefits from punishing another. If the population initially consists of many policers, honest-types who do not punish will invade and this will in turn result in the population being swiftly taken to the all cheat equilibrium. Honesty cannot be sustained.

It is tempting to think that any incentive structure can yield these favorable results, as long as policers are rewarded for catching cheaters and cheaters are punished when caught. This, unfortunately, is not always the case. Certain parameter values result in undesirable behavior, such as the community cycling in an extreme fashion from (nearly) all being honest to (nearly) all cheating to (nearly) all policing. Figure 2 (a) illustrates this case. Under other parameter values the population converges to a polymorphic equilibrium in which the population is split between policers and cheats (Fig. 2 (b)).[13]

Figure 3 depicts which of the above dynamics we'd expect in various regions of parameter space (see the Appendix for a description of how this figure was created). We see that under a wide variety of parameter values the resulting dynamic is that of Fig. 1(a) – the population spirals towards a state in which all three strategies are present but the overwhelming majority of the agents play either "Honest" or "Police." This arrangement can obviously only be sustained if p, the cost of being caught as a cheater, is greater than θ, the benefit of cheating. Additionally, it is worth noting that b, the benefit associated with successfully catching a cheat, must be greater than p. In other words, only incentive structures in which the reward for catching a cheat is greater than the cost of being exposed as a cheat can sustain peer-policing in science. Whether or not the incentive structure scientists face in the real world satisfies the above conditions is an open question.

4. DISCUSSION

I have explored one means by which a community of truth-seekers can minimize the amount of misbehavior that takes place in science: peer-policing. Peer-policing rests on the assumption that scientists are not angels, but boundedly rational agents pursuing their own self-interest. As a result, scientists must somehow be incentivized to police others. In other words, punishment in science must be opportunistic and not altruistic. This was Hull's fundamental insight, and I believe the results of this paper support his optimism that science *can* function as a "self-policing system." For a wide variety of parameter values, a decentralized community of scientists can minimize fraud, fabrication and corner cutting if, among other things, scientists can benefit by investigating the claims of

11 In all of these cases the polymorphic equilibrium is an attractor. The population spirals toward the equilibrium from any interior point.

12 This is formally equivalent to the prisoner's dilemma with costly punishment.

13 There are also cases where this polymorphic equilibrium is unstable, in which the population heads to the cheater equilibrium in the long run.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

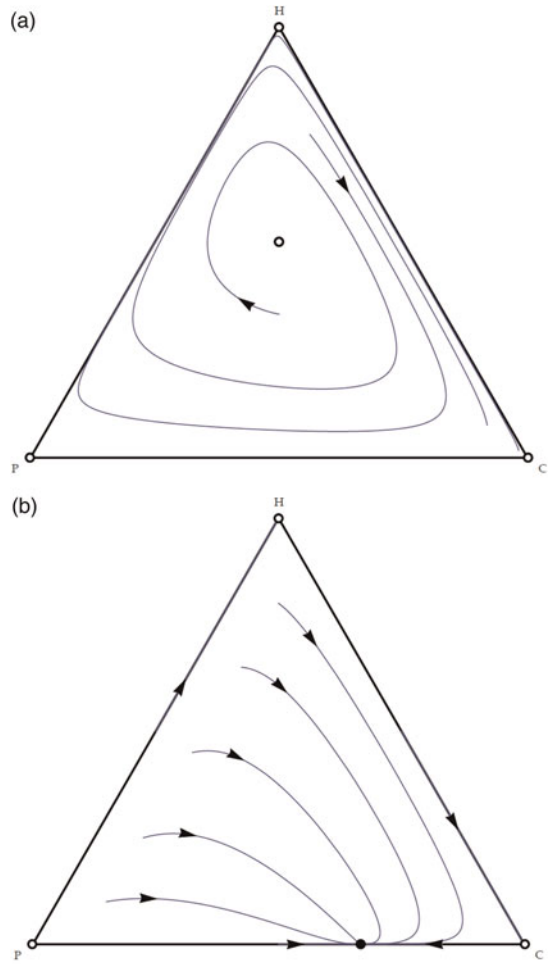


Fig. 2. (colour online) (a) Dynamics of the Merit-Demerit game when $\mu = 10$, $\theta = 1$, $p = 4$, $b = 2$, $c = 0.5$. The polymorphic equilibrium is not stable and we get spiraling to the border of the simplex and (b) The Merit-Demerit game when $\mu = 10$, $\theta = 1$, $p = 4$, $b = 6$, $c = 0.5$. All interior points of the simplex head to the polymorphic equilibrium consisting of only policers and cheaters.

others. Interestingly, the chief assumption of the economic approach to science – that scientists are self-interested – is the root of both the problem and solution. Misbehavior exists since scientists care about maximizing their reputation, but the temptation to fabricate data can be kept in check if exposing fraud and corner-cutting is itself reputation-enhancing.

Throughout this paper I've purposely been conflating two distinct issues, fraud and error. There is, of course, a world of difference between these two. Fraud is intentional. It is a deliberate and calculated attempt to fool the scientific community. Errors, however, are commonplace in science and, at their very worst, a result of careless and sloppy work. I believe my model can address both the problem of fraud and mistakes. It is obvious how the Merit-Demerit game of Table 1 maps onto the case of fraud. It is less clear what the

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

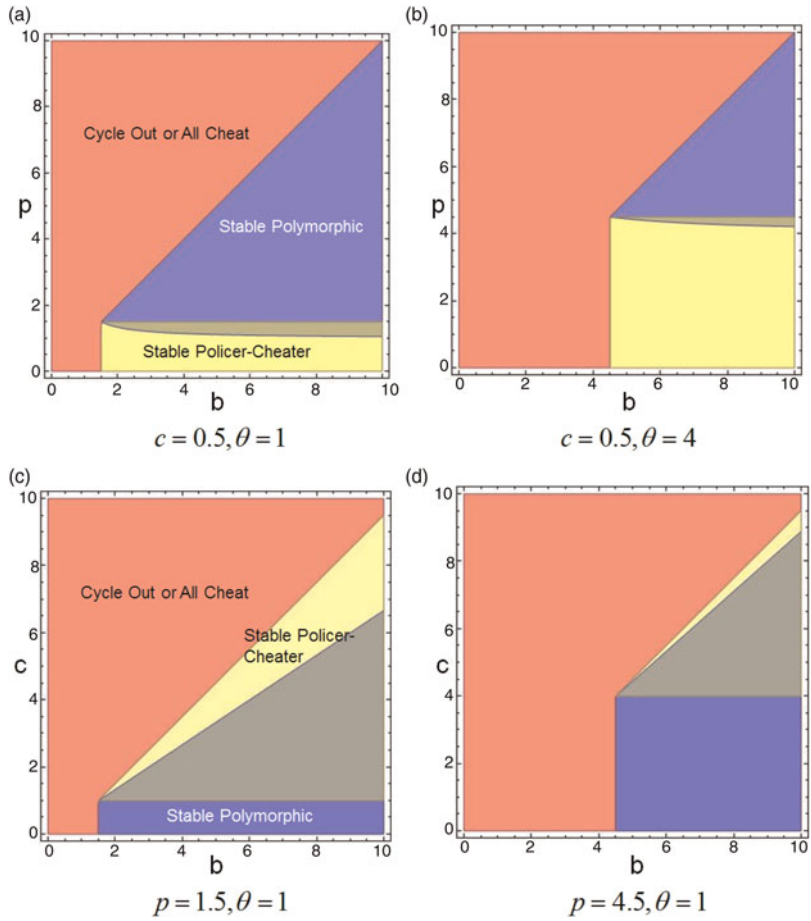


Fig. 3. (colour online) Projection of the five-dimensional parameter space. μ is held constant at 10 in all figures. Parameters *c* and *b* denote the cost and benefits associated with policing, *p* and θ the costs and benefits of cheating. Regions labeled "Cycle Out or All Cheat" denote areas of parameter space where the behavior of the community either heads to the All-Cheat equilibrium or cycles outward to the border of the simplex. Regions labeled "Stable Polymorphic" denote parameters that support a lone polymorphic equilibrium in which all three strategies are present. Regions labeled "Stable Policer-Cheater" denote regions where the community heads to a polymorphism where only Cheats and Policers exist. The gray regions indicate parameters that sustain both types of polymorphic equilibria.

relationship between the Merit-Demerit game and error is. If we interpret the "cheat/cut-corners" strategy as "do shoddy, unreliable work," then the scientist presumably has more time to spend on other endeavors. He is thus able to publish more in a given year than his more fastidious peers. This in turn accounts for his enhanced payoff of μ + θ.

Of course this is not to say that there aren't important differences between fraud and accidental error in science. Accusations of fraud are treated much differently than those of error. In particular, it is too often the case that the careers of the "whistle blower" are put in jeopardy. Hull makes this exact point, and references the sad case of Margaret O'Toole, who was involved in the Baltimore–Imanishi-Kari scandal (Hull 2000). In the real world it may often be the case that those who attempt to uncover foul-play in science are not

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

rewarded, but instead penalized. Another difference between fraud and error is that the penalty associated with being caught as a fraud is exceptionally high. Most cases of fraud result in the guilty party leaving academia altogether. Yet as we noted in the past section, for peer-policing to be feasible the benefits of policing must exceed the cost of being caught as a cheat. This seems to suggest that peer-policing may be quite effective at ensuring scientist's do not engage in careless or sloppy work, but hard-pressed when dealing with fraudulent behavior.[14]

The realization that peer policing is ill-suited to deal with fraud is just one of the normative implications of my analysis. My study of peer-policing indicates that for the population to consist primarily of honest types both the benefits of catching a cheat and the cost associated with being caught should be rather high (with the former exceeding the value of the latter).[15] Yet how are individuals "punished" for cutting corners? Surely publishing an article riddled with errors isn't enough to be expelled from academia. Punishment, I argue, is administered when the community at large forms a negative opinion of the culprit, which in turn requires that the culprit's shoddy work be publicly exposed as such. Yet communicating to the community at large that a research article is flawed is a more difficult task than one would expect. Even when an article is retracted due to error or misconduct the community may not take note and continue to cite the work.[16] Thus in order to ensure that the community is informed of the erroneous article, journals should do more to publicize retractions and in addition make clearer the reason the article was retracted.[17] This would have the effect of increasing the value of p (punishment meted out to the cheater) in my model.

My model also implies that in order to minimize the amount of misbehavior in the community more value should be placed on the replication of past experiments. As it currently stands, the scientific community in general desires and rewards novel results and pays little attention to replication (especially if the results of the replication are consistent with the results of the original experiment). If attitudes toward replication changed then more scientists would be willing to devote time and energy toward replicating past experiments. In terms of my model, this would correspond to a decrease in the parameter c (the costs associated with policing).[18] If replications were appreciated and frequently published, then even if a policer failed to "catch" a cheater by replicating her experiment, the policer would still benefit and perhaps even be able to publish the replication.

As is the case with all formal work, I make use of a number of idealizing assumptions. While these assumptions simplify the analysis, they by no means undermine the value of my model. In fact, the cycling in Fig. 1(a) can be found in other formal work exploring honest behavior and corruption. Bicchieri and Duffy (1997), for example, develop a model in which society continually vacillates between honest and corrupt behavior. Further work aimed at relaxing stringent assumptions is required in order to gain a fuller understanding of peer-policing in science.

14 While David Hull is a firm believer that science is a self-policing system with respect to error, he is skeptical that peer-policing can deal with fraud (Hull 1996b).

15 Recall that the proportion of honest types in the polymorphic equilibriums is given by $1 - \theta/p - c/b$.

16 Retracted articles are, in fact, heavily cited in certain cases. See Pfeifer and Snodgrass (1990) and Budd et al. (1998) for a survey of retracted articles in medicine.

17 While it would be nice if journals did more to publicize retractions, journals have little incentive to do this. After all, one does not want to acquire the reputation of being the Journal of Retractions.

18 Recall that the proportion of cheaters in the polymorphic equilibrium is given by c/b. Thus lowering c reduces the number of cheats in the population.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

Lastly, I'd like to emphasize that the topic of misbehavior has been for the most part completely overlooked in the social epistemology of science. Yet we know that from time to time data *is* fabricated and that scientists sometimes *don't* take enough care with their research. Thus any account of why science is the success that it is must address how these communities of truth-seekers can successfully curtail fraud and corner-cutting. This paper is a modest first step in this direction, but further inroads must be made to explore the full range of options we have at our disposal.[19]

5. TECHNICAL APPENDIX

This technical appendix briefly outlines the process by which we construct Fig. 3. We begin by first calculating the average payoff to each of the three strategies in the Merit-Demerit game.

$$EU(H) = \mu$$

$$EU(P) = \mu - c + (1 - x_1 - x_2)b$$

$$EU(C) = \mu + \theta - px_2$$

Where x_1 and x_2 represent the proportion of honest-types and policer-types in the population. We can calculate the average payoff of the population simply enough:

$$\mu + \theta - \theta x_1 + (b - c - p - \theta)x_2 + (p - b)x_1x_2 + (p - b)x_2^2$$

Plugging into the replicator dynamics, we get the following:

$$\dot{x}_1 = x_1(-\theta + \theta x_1 - (b - c - p - \theta)x_2 - (p - b)x_1x_2 - (p - b)x_2^2)$$

$$\dot{x}_2 = x_2(b - c - \theta + (\theta - b)x_1 + (c + p + \theta - 2b)x_2 - (p - b)x_1x_2 - (p - b)x_2^2)$$

To assess the stability of the rest points, we first construct the Jacobian matrix.

$$j = \begin{bmatrix} \frac{\partial \dot{x}_1}{\partial x_1} & \frac{\partial \dot{x}_1}{\partial x_2} \\ \frac{\partial \dot{x}_2}{\partial x_1} & \frac{\partial \dot{x}_2}{\partial x_2} \end{bmatrix}$$

19 I'd like to thank Simon Huttegger, Hannah Rubin, Gregory McWhirter, Brian Skyrms, Kyle Stanford, Kevin Zollman, Rainer Hegselmann and Sarah Lawsky for helpful discussions and comments. I'd also like to thank the participants of Choosing the Future of Science: the Social Organization of Scientific Inquiry at the Center for Philosophy of Science, University of Pittsburgh, the participants of the Works in Progress series at the University of California, Irvine, the members of the Equity Reading Group and the helpful comments of an anonymous referee.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

Where,

$$\frac{\partial \dot{x}_1}{\partial x_1} = -\theta + 2\theta x_1 - (b - c - p - \theta)x_2 - 2(p - b)x_1 x_2 - (p - b)x_2^2$$

$$\frac{\partial \dot{x}_1}{\partial x_2} = -(b - c - p - \theta)x_1 - (p - b)x_1^2 - 2(p - b)x_1 x_2$$

$$\frac{\partial \dot{x}_2}{\partial x_1} = (\theta - b)x_2 - (p - b)x_2^2$$

$$\frac{\partial \dot{x}_2}{\partial x_2} = (b - c - \theta) + (\theta - b)x_1 + 2(c + p + \theta - 2b)x_2 - 2(p - b)x_1 x_2 - 3(p - b)x_2^2$$

Using *Mathematica* we find the eigenvalue of the above Jacobian. To ascertain the stability properties of the polymorphic equilibrium, we plug in $1 - \frac{\theta}{p} - \frac{c}{b}$ for x_1 and $\frac{\theta}{p}$ for x_2. This yields two identical eigenvalues:

$$\frac{1}{2b^2p^2}(-b^2cpt + bcp^2t - b\sqrt{c}p\sqrt{t}\sqrt{-4b^2p^2 + 4bcp^2 + b^2ct + 4b^2pt - 2bcpt + cp^2t}).$$

The polymorphic equilibrium is an attractor when the real part of the above eigenvalue is negative. The stability properties of the other rest points are ascertained in a similar fashion.

REFERENCES

Bicchieri, C. and Duffy, J. 1997. 'Corruption Cycles.' *Political Studies*, 45: 477–95.

Boyd, R. and Richerson, P. 1992. 'Punishment allows the Evolution of Cooperation (or anything else) in Sizable Groups.' *Ethology and Sociobiology*, 3: 171–95.

Boyd, R., Gintis, H., Bowles, S. and Richerson, P. 2003. 'The Evolution of Altruistic Punishment.' *Proceedings of the National Academy of Sciences USA*, 100: 3531–5.

Budd, J., Sievert, M. and Schultz, T. 1998. 'Phenomena of Retraction: Reasons for Retraction and Citations to Publications.' *Journal of the American Medical Association* 280: 296–97.

Cant, M. and Johnstone, R. 2006. 'Self-serving Punishment and the Evolution of Cooperation.' *Journal of Evolutionary Biology*, 19: 1383–5.

Colman, A. 2006. 'The Puzzle of Cooperation.' *Nature*, 440: 744–5.

Cox, J. and Goldman, A. 1994. 'Accuracy in Journalism: An Economic Approach.' In F. Schmitt (ed.), *Socializing Epistemology: The Social Dimensions of Knowledge*. London: Rowman and Litterfield Publishers.

Frang, F., Steen, R. and Casadevall, A. 2012. Misconduct accounts for the majority of retracted scientific publications. *Proceedings of the National Academy of Sciences USA*, 109: 17028–33.

Hull, D. 1996a. 'Why Scientists Behave Scientifically.' *MRS Bulletin*, 21: 72.

——. 1996b. 'What's Wrong with Invisible-hand Explanations?' *Philosophy of Science*, 64: S117–25.

——. 2000. *Science and Selection: Essays on Biological Evolution and the Philosophy of Science*. Cambridge: Cambridge University Press.

Kitcher, P. 1990. 'The Division of Cognitive Labor.' *Journal of Philosophy*, 87: 5–22.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press

Nakao, H. and Machery, E. 2012. 'The Evolution of Punishment.' *Biology and Philosophy*, 27: 833–50.

Panchanathan, K. and Boyd, R. 2004. Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature* 432: 499–502.

Pfeifer, M. and Snodgrass, G. 1990. The continued use of retracted, invalid scientific literature. *Journal of the American Medical Association* 263: 1420–3.

Richerson, P. and Boyd, R. 2005. *Not By Genes Alone: How Culture Transformed Human Evolution*. Chicago, IL: University of Chicago Press.

Rohwer, Y. 2006. 'Hierarchy Maintenance, Coalition Formation, and the Origins of Altruistic Punishment.' *Philosophy of Science*, 74: 802–12.

Sandholm, B. 2010. *Evolutionary Games and Population Dynamics*. Oxford: Oxford University Press.

Schlag, K. 1997. Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78: 130–56.

Sripada, C. 2005. Punishment and the strategic structure of moral systems. *Biology and Philosophy*, 20: 767–89.

Strevens, M. 2001. 'The Role of the Priority Rule in Science.' *Journal of Philosophy*, 100: 55–79.

Taylor, P. and Jonker, L. 1979. 'Evolutionary Stable Strategies and Game Dynamics.' *Mathematical Biosciences*, 40: 145–56.

Wible, J. 1992. 'Fraud in Science: An Economic Approach.' *Philosophy of the Social Sciences*, 22: 5–27.

JUSTIN BRUNER is a PhD student in the Department of Logic and Philosophy of Science at the University of California Irvine. His main research interests are in game and decision theory, social contract theory and the philosophy of biology.

https://doi.org/10.1017/epi.2013.34 Published online by Cambridge University Press