

Why Athletic Doping Should Be Banned

ERIC CHWANG

ABSTRACT *So long as a ban is enforceable, large private athletic institutions—such as Major League Baseball and the National Collegiate Athletic Association—should not allow their athletes to take performance-enhancing drugs. The argument I present is game-theoretic: though each athlete prefers unilateral permission to dope over a universal ban, he also prefers a universal ban over universal permission to dope. That is because, while doping improves absolute measures of performance, it does not improve relative performance if many athletes dope. Large private athletic institutions should honour their athletes' preferences and should not enact any policy that gives only some athletes but not others permission to dope. Thus, they should ban doping. My paper examines and defends this game-theoretic argument. After explaining the argument, I compare it (favourably) to other arguments for the same conclusion. I then discuss whether the game-theoretic argument counterintuitively fails to justify banning some forms of doping (no), and whether it counterintuitively justifies banning things other than doping (also no). After arguing that a doping ban on game-theoretic grounds is neither wrongly paternalistic or nor wrongfully coercive, I end by discussing some limitations of the argument.*

1. Introduction

So long as a ban is enforceable, large private athletic institutions—such as Major League Baseball (hereafter 'MLB') and the National Collegiate Athletic Association (hereafter 'NCAA')—should not allow their athletes to take performance-enhancing drugs. Let us abbreviate 'take performance-enhancing drugs' as 'dope' and 'large private athletic institutions' as 'we', taking care to note that in the context of this paper 'we' will never refer to government entities. Let us also leave implicit the assumption that a ban is enforceable: the scientific issue of testing feasibility for all possible doping agents is enormously complex and rapidly changing, and I will simply ignore it.¹ Using these abbreviations, my thesis is that *we should ban athletic doping*.

The question I address differs in important ways from all of the following: (a) whether athletes should dope, (b) whether athletic doping should be illegal, (c) whether other forms of (artificial) athletic performance enhancement besides doping—e.g. high performance swimsuits or genetic modification—should be banned (by private bodies or the government), and (d) whether doping should be banned outside regulated athletics. I will mention some of these issues tangentially, and I come back to some of them at the end of this paper, but they are not my main focus.

The rationale I present and defend is not new.² Still, the argument needs elaboration, as we will see, and the extant literature on doping has conflated its essential idea with simpler and closely related yet inadequate ideas. That is bad, because the argument I develop is an improvement over other arguments for the same conclusion, and it evades various recent objections from critics who think doping should be permitted.³

The paper is divided into six sections, including this one. In the next section, I explain the game-theoretic reason we should ban doping and compare it to other arguments for the same conclusion. In Section 3, I discuss whether this game-theoretic argument counterintuitively fails to justify banning some forms of doping, and in Section 4 I discuss whether it counterintuitively justifies banning things other than doping. In Section 5, I ask whether a doping ban is wrongly paternalistic or coercive. I discuss some limitations of the argument in Section 6.

2. A Game-Theoretic Reason to Ban Athletic Doping

2.1. *The Basic Argument*

Each individual athlete wants a competitive advantage over his rivals, often even to the point of doping. However, unlike food or most over-the-counter drugs such as vitamins, doping is harmful to those who dope, mainly because of the physiological side effects, but also because of the opportunity cost of time and resources spent acquiring dope. Thus, athletes prefer not to dope if they do not have to. Still, doping improves performance, so athletes prefer the option of self-harm plus improved performance—i.e. doping—to required abstention. The best possible scenario for any individual athlete is where he alone is permitted to dope, whereas the worst possible scenario is where he alone is required to abstain. Thus, where being higher in the ranking indicates being preferred, the individual preference ranking for any rational (actual) athlete is (should be):

Doping permitted for me alone.
 Doping permitted for all.
 Abstention required for me alone.

If every athlete had this preference ranking, permissive doping might well lead to the tragic situation where everyone dopes, but with no competitive advantage for anyone.

The key to escaping this tragedy is to highlight the neglected possibility that doping be banned. Inserting that possibility to our athlete's preference ranking, we have:

Doping permitted for me alone.
Abstention required for all.
 Doping permitted for all.
 Abstention required for me alone.

Every rational (actual) athlete will (should) prefer a total doping ban to completely permissive doping, because (a) completely permissive doping will not give that athlete an advantage over his competitors (I will discuss the possibility that doping is differentially effective in Section 3), and (b) as long as you cannot gain any competitive advantage, you might as well avoid self-harm. We should protect athletes from harm, and we would like to give them what they want. We cannot implement each athlete's most favoured policy of unilateral doping for that particular athlete and no one else. Therefore, we choose the universally next most favoured option of mandatory abstention for all. In other words, we should ban athletic doping.

How to label this argument? It can be called *strategic*, because an individual athlete is in a sense willing to give up the option of doping, so long as everyone else gives up that

option too: a complete ban on doping is preferable (for all athletes) to complete permission. For the same reason, the two-competitor version of the argument can be modelled as a *prisoner's dilemma*, where doping is defection, and abstention is cooperation. The multi-competitor version can be modelled as a *tragedy of the commons*, where doping is over-grazing, and abstention is sensible grazing. Indeed, in the classic tragedy of the commons for public goods such as for clean air, regulatory jurisdiction is at best controversial, and there is no easy analogue to athletic governing bodies. Thus, pollution, modelled as a tragedy of the commons, presents a difficult problem. By comparison, doping in competitive sports governed by organizations such as MLB is a trivial problem, precisely because of clear governance.

One might also describe the argument as *harm-based*, because our concern is with protecting athletes from harm. Indeed, one might even call it *paternalistic*, a possibility I will examine in more detail and eventually reject in Section 5. For now, note that appealing to (protection from) harm does not fully capture the force of the argument. Yes, one key assumption in the argument is that doping harms those athletes who dope, but another crucial assumption is that if sufficiently many athletes dope then no competitive advantage is gained, while all dopers are harmed. This sort of scenario is analogous to an *arms race*, where countries spend an ever-increasing sum on the military (thereby forgoing the opportunity to allocate those funds to other areas, such as domestic improvements), but without gaining competitive advantage over rival countries that are likewise escalating military spending. More generally, these sorts of situations can be described as *positional treadmills*, where a competitive agent expends valuable resources without gaining any advantage for her efforts, because all her competitors are doing the same.⁴

To show the importance of the positional treadmill assumption, it will be useful to pause and contrast athletic doping to, say, musical doping. Suppose a musician takes a prescription-only drug, such as a beta-blocker, to calm herself before a big performance, thereby enhancing that performance.⁵ Although the quality of her performance relative to that of other musicians may partly determine what future jobs she is offered in a competitive environment, we also care about the quality of her performance for its own sake. (Here 'we' refers to listeners and performers.) That is, we care about the absolute quality of her performance rather than merely about it being relatively superior to that of rival musicians. That is much less true in competitive sports, where we (fans and athletes) mainly care about relative quality—who is better, which team won, etc.—not absolute quality. More cautiously, we mainly care about relative quality in competitive sports where the level of competition is already very high, as it is in MLB and even any NCAA division. Thus, athletic performance enhancement in competitive sport (where the level of competition is antecedently high) results in a purer positional treadmill than musical performance enhancement does, and my argument correspondingly applies much more strongly to athletic doping than to musical doping.

Given all these ways of labelling the argument, let us just use a loose umbrella term and call it *game-theoretic*. We are now in a position to make the form of this argument more explicit, though of course the premises involved are still imprecise, and the argument form is still rough.

1. Doping harms dopers. [Factual Harm Premise]
2. While doping may improve absolute measures of performance, it provides no competitive advantage if others also dope. [Factual Treadmill Premise]

3. Therefore, athletes prefer a total doping ban to permissive doping for all. [From 1 and 2]
4. We should not enact any policy that gives some athletes permission to dope but not others. [Moral Impartiality Premise]
5. Within some constraints, we should give athletes what they want. [Moral Beneficence Premise]
6. Therefore, we should ban doping. [3, 4, 5]

2.2. *Rough Patches*

In Sections 3 and 4 I will examine some of the assumptions of the game-theoretic argument in more detail. For example, in Section 3 I will discuss the possibility of non-harmful dope (contrary to premise 1) and the possibility that some competitors may benefit more than others from permissive doping (contrary to premise 2). In this sub-section, I will point out parts of the argument that are a bit rough, though not dangerously so.

One way in which this argument is rough is that I have not proved that the inference from steps 1 and 2 to 3 is valid. The formalization and modelling required for this would be bothersome and unhelpful, so I prefer to leave the argument informal.

Relatedly, I have suppressed the assumption that (spectators and) athletes care a great deal about beating their competitors, and not just about the quality of their performance in absolute terms. Thus, for example, the analogue of premise 3 would not follow from harm and treadmill premises for musicians taking beta-blockers. Let us call sports where the athletes and spectators care mostly about beating rivals *competitive*, using *non-competitive* to refer to those sports ('activities', for those who think competition is an essential element of sport) where we care mostly about the absolute quality of performance. The game-theoretic argument applies only to competitive sports (with an antecedently high absolute level of performance), not to non-competitive ones, and of course we may debate which real world sports should be thought of as competitive, by this definition.

Another way in which the argument is rough is that it finesses the move from what rational athletes desire to what we should do for actual athletes. Relatedly, I will slide back and forth between talking about *preferences* of athletes (which may be irrational) and *harms* to them (which athletes suffer even if, irrationally, they do not care). The argument may be vulnerable here, but I will not pause to discuss ways to shore it up.

Also, I have not spelled out the relevant constraints mentioned in step 5. Athletes want a lot of money, for example, but that does not mean we are obligated to indulge them without limit. I will invoke a particular constraint—what spectators want—in Section 4, but otherwise I will not bother to discuss other constraints, because the argument does not run afoul of any. Indeed, a plausible constraint on satisfying preferences is that we should strive to satisfy only those preferences that a prudentially rational athlete would have, in which case the problem of the previous paragraph is eliminated.

Finally, I have focused on the consequences of doping only for those athletes who dope. Thus, for example, I ignore the effect of permissive doping on children, families, and even society more broadly, perhaps via reputation effects, the economy, or the environment. Of course these effects matter too, but athletic organizations have a primary fiduciary responsibility to protect their athletes that goes beyond their responsibilities to others.

2.3. Other Arguments

The game-theoretic argument for a doping ban is better than other arguments for the same conclusion. It appeals to concrete harms to athletes and their (rational) preferences to avoid those harms. This puts the game-theoretic argument on more solid footing than arguments that claim that doping is *unnatural* and arguments that claim that achievements accomplished via doping are *inauthentic*, where inauthenticity may stem from the achievement not being accurately described as the athlete's or from its coming too easily.⁶ Appeals to naturalness and authenticity as virtues are, at the very least, controversial. Even if one could craft those distinctions precisely enough to cut at just the right joints, showing that doping is unnatural, or that doped achievements are inauthentic, in the relevant senses, while simultaneously showing that other forms of enhancement are natural, and other forms of achievement authentic, why should we care about naturalness or authenticity to begin with? The game-theoretic argument, in contrast, appeals to a consideration everyone agrees is important: the desire to avoid harm.

The game-theoretic argument emphasizes harm to *athletes*, which puts it on more solid footing than arguments that claim that doping harms *sport*.⁷ Now, the 'harm to sport' label is slippery, because it can be used in a catch-all way to refer to many different arguments, for example those of the previous paragraph. Perhaps a catchy way to express the reason why we care about naturalness or authenticity is that when these values are diminished, sport is harmed. In these cases the criticisms in the prior paragraph still apply: these alleged values are controversial, whereas the desire to avoid harm is not.

The 'harm to sport' label can also be used to suggest that we should care about some thing—namely sport—over and above its participants and spectators (and perhaps other sentient creatures it affects). For example, if the National Football League (hereafter 'NFL') allowed its players to use (safe) jet packs on the field, some might argue that this would harm the sport of (American) football, regardless of its effect on players or fans. The ontological presupposition here is dubious, but the alleged value at stake is still a lot more controversial than the desire to avoid harm. Further, the alleged value is not even very strong. Thus, even a traditionalist who wants to ban jet packs from the NFL may think it reasonable, or at least morally permissible, for a rival league to introduce jetpacks as a way of differentiating itself from the NFL. He may cringe at the thought, because of the harm to (American) football, and he may object to its still being labelled '(American) football', but this is not the same as to think that the rival league should be forbidden from introducing jet packs. In contrast, introducing dope is not a morally legitimate way for a rival football league to differentiate itself from the NFL, because the argument for a doping ban applies both to the NFL and to all its competitors.

Relatedly, some have defended a doping ban on grounds that *public support* for sports would falter if athletes were allowed to dope.⁸ This is at best speculative. The game-theoretic argument is not speculative in this way. Relatedly, one might argue that athletic doping should be banned because professional athletes are *role models*.⁹ Here the intended suggestion—a complete ban—seems out of proportion with the alleged harm. After all, children also see athletes drink, smoke, and philander, yet we do not think athletes should be banned from engaging in these activities just because they are serving as poor role models by doing so. In contrast, the positional treadmill problem is significant enough to be proportional to a complete ban.

The game-theoretic argument is more sophisticated than the simple harm-based argument that doping should be banned because it harms athletes. As with the role model argument, the suggestion seems out of proportion with the alleged offense. Even if smoking is harmful and (especially) irrational for athletes, still, that fact alone seems not to warrant a smoking ban for athletes. As I will discuss in more detail in Section 5, the game-theoretic argument, with its essential factual treadmill premise (2), is not wrongfully paternalistic in the way that a simple harm-based argument may be.

Finally, an obvious argumentative strategy for a doping ban is that allowing dope would be *unfair*. However, these arguments must rely on a controversial view of fairness, for it seems that fairness is preserved if everyone is permitted to dope. Indeed, the impartiality premise of my argument (premise 4) is compatible with that thought. At any rate, it is not immediately obvious how permissive doping is any more a culprit against fairness than, say, differential wealth between rival countries in the Olympics.¹⁰ Thus, the game-theoretic argument can motivate a ban on performance-enhancing drugs without simultaneously motivating (or perhaps ‘without simultaneously needing *ad hoc* epicycles to avoid . . .’) rigid salary caps or equal funding for Olympic rivals or other controversial applications of stronger egalitarian principles. In this sense, the scope of the game-theoretic argument is narrower and more specific than that of fairness based arguments: it can justify a doping ban without thereby also justifying many more radical changes to competitive sports.

3. Proves Too Little?

Perhaps the game-theoretic argument does not go far enough to rule out all forms of doping that, intuitively, we think should be banned. Recall that two descriptive facts about doping are essential to the game-theoretic argument. First, performance-enhancing drugs have harmful side-effects to those who take them (the factual harm premise), and second if sufficiently many athletes dope then there is minimal relative competitive gain (the factual treadmill premise). The game-theoretic argument does not justify a ban in doping situations where either aspect is missing. I will consider these two possibilities in order.

3.1. *Non-Harmful Drugs*

First, the game-theoretic argument does not justify a ban on performance-enhancing drugs that lack harmful side-effects. Are there any such drugs? Clearly some substances, when ingested in dosages that help performance, do little to no harm to athletes. Food and water are examples. Hyponatræmia is a dangerous consequence of water intoxication, but the incredibly high levels of water ingestion needed to cause hyponatræmia (in an otherwise healthy adult) do not help athletic performance. Unhealthy foods are harmful in the long run even while helping performance in the short run, but they do not improve performance beyond the gains of eating healthy food. Thus, an athlete who (knowingly) chooses unhealthy food over healthy food cannot be doing so for the sake of improving athletic performance. Also, remember that athletes care about more than mere physiological side effects. They also care about opportunity costs; that is, they care about how much time, effort, and money they have to invest in order to acquire their

dope. The game-theoretic argument still goes through if these costs are sufficiently severe.

In any case, prescription drugs arguably all have harmful physiological side-effects: that is why they are regulated, after all. Consider, as an example, erythropoietin ('EPO'). EPO is a naturally occurring hormone that stimulates red blood cell production. Exogenous EPO can be an essential part of the treatment of certain forms of anemia (insufficient red blood cells). It has also been used as a doping agent in sports such as cycling and running, where differential cardiovascular efficiency is a key competitive differentiator between winners and losers. EPO, whether endogenous or exogenous, is dangerous in large concentrations (it can increase the chance of adverse clotting, leading to heart attacks or strokes, for example), but it can be used safely.¹¹ Indeed, one way to increase the natural concentration of EPO in one's blood is just to train at high altitudes, where there is less oxygen. Now, although exogenous EPO can be administered safely in low doses, if we allow EPO doping then athletes will end up in an arms race scenario where they will be tempted to use ever-increasing concentrations of it just to stay competitive, even past the point where it becomes dangerous. (This assumes, as is empirically well-founded, that increasing concentrations of EPO continue to improve performance well past the point at which it becomes dangerous.) Thus, what looks initially like a promising candidate for a non-harmful drug turns out to be harmful after all.

Steroid doping is similar. Anabolic steroids occur naturally and can be regulated by exercise, and exogenous steroids are safe in small enough doses. Still, if we allow steroid doping then an ever-escalating arms race is likely to develop, in which case they become dangerous.¹² The theoretical point remains, however: my argument does not motivate banning performance-enhancing drugs that are safe (and cheap). I view this specificity as a positive feature of the argument, however, not a flaw. That is because my goal is not first to determine what all the performance-enhancing drugs are, and then to find a rationale to ban them all. Rather, my goal is more akin to rational reconstruction as practiced by the logical positivists: I am trying to find a strong rationale for a something approximating a popular practice (a doping ban) that many antecedently find plausible, and in those few instances of the popular practice where the rationale does not apply, to remain silent.

Still, at this point the game-theoretic argument faces a serious objection: if dope is safe in small doses but harmful at large doses, why not ban large doses but permit small ones? The response to this relies on three inter-related points. First, when I said earlier that there are safe doses of dope I was simplifying: there is no binary cut-off between safe and unsafe doses of performance-enhancing drugs. Rather, there is a continuum of safety, between more and less safe doses. Put probabilistically, any increase in dosage leads to a corresponding increase in probability of harm. Second, even if there is a sharp cut-off between safe and unsafe concentrations of dope, that cut-off is both subjective and controversial. It is subjective in that it differs from person to person (and from situation to situation for any given person), and it is controversial in that even when one knows all the descriptive facts about a particular case one might still dispute the normative threshold between safe and unsafe. Finally, given that we want to restrict doping at least somewhat, the marginal cost of moving to a total ban is small. That is, once we agree to incur the cost to test athletes to see how much they have been doping, the additional incurred cost for setting our tolerance level at 0, rather than at some

intermediate level, is negligible. Couple all these points with the positional treadmill fact that athletes do not gain competitive advantage even (especially) from very small doses of dope, and the best policy might still be an absolute ban rather than a cap on permissible doping.

These responses are instructive not only in those cases where they succeed but also in cases where they fail. For example, ingredients in routinely available generic medicines that alleviate mild ailments such as the common cold might also enhance performance, perhaps because they contain mild stimulants, which may help ramp up the immune system. Presumably the utility gain from moving away from an absolute ban in this case *is* worth the cost of haggling over where to set a cap on permissible dosages of cold medicines: we want to allow athletes to be able to get over their colds as soon as they can. In other words, perhaps we should allow some use of such medicines, even though doing so will be costly, in terms of squabbling about the appropriate, non-zero dosage cut-off.

3.2. *Competitive Advantage*

The simple game-theoretic argument does not condone a doping ban in situations where doping differentially improves the performance of some athletes more than it improves the performance of others. That is because in these situations, some athletes, such as the one who stands to gain the most from doping, may rationally prefer complete permissiveness to a complete ban. What to say about this possibility?

One sort of scenario we can dismiss immediately is that some type of player in a team sport will benefit more than rival types in the same sport. For example, suppose that permission to use steroids in major league baseball would increase batting averages, home run totals, and runs per game. Even so, the game-theoretic argument can still justify a ban on steroids. That is because, although hitters may benefit from steroids more than pitchers do, individual hitters still need not benefit more than other hitters do. Thus, although offensive statistics would inflate, win-loss percentage may still be unaffected. And, of course, the competition amongst rival hitters for the batting crown, home run title, and other offensive accolades will still be just as fierce, even if the relevant absolute numbers are all inflated.

Another rationale we can likewise dismiss is that all current players prefer permissive doping because that will help them break the records of former players. In this respect, doping is analogous to many other improvements over time, for example in technology, technique, or strategy, that help current athletes gain an advantage over former ones. The problem with this hypothesis about athletic preferences is not that it is false. Rather, the problem is that athletic organizations should look out for the interests of their former athletes as well as of their current ones. In other words, athletic organizations should not favour current athletes at the expense of past athletes.

More dangerous to the game-theoretic argument is the straightforward possibility that some athletes will benefit more from doping than other (current) athletes (of the same type) do. However, the game-theoretic argument still goes through even in this case. That is because, even if dope in fact differentially benefits one athlete more than another, there is typically no good reason for any individual athlete to believe that dope will increase his performance so much more than it does that of his competitors that the probability of his winning any prize will also increase enough to outweigh its harms. In other words, the

factual treadmill premise (#2) from Section 2 can be replaced by an epistemic variant—namely, no athlete *has good reason to think he* will acquire a relative advantage over his competitors if everyone is allowed to dope—and the rest of the argument still goes through without change.

Further, even if some athlete does have good reason to think dope will differentially improve his performance more than that of his competitors, a total ban is still justified. Before I explain why this is so, let me emphasize that I think this logical possibility is in fact very improbable. Do not take the relative length of my explanation to follow to correlate with my estimate of the plausibility of our starting assumption, that some athlete may have special epistemic access to the differential effects of dope on him versus his competitors. The primary response to this hypothesis is to deny it.

But in what follows let us suppose, for the sake of argument, that the differential effects of dope on different athletes are common knowledge to all competitors. It will be useful to divide the argument here into cases, first where there are many more losers than winners, and second where we loosen the concept of winning sufficiently so that there can be many winners.

First, suppose that there are many more losers than winners. For example, suppose the fastest three runners of a hundred get a prize, and they all benefit to a different degree from a fixed dose of EPO. For simplicity, suppose that it is a complete matter of chance how the runners would place against each other in a non-doped race, but that their relative placements in an EPO-permissive race are completely determined. As specified earlier, we also suppose that all these details are common knowledge among the runners. In this case, we can expect at most three of the hundred to prefer permissive doping to a ban. (I say ‘at most’ because some of the winners may still think that their prize is not worth the cost of doping.)

However, this set of preferences still justifies a complete doping ban. That is because we should implement the policy that the vast majority of athletes want implemented, especially when the contrary position harms some of the athletes; unanimity is unnecessary for the argument. In other words, to deal with this case I can just retract and qualify the statement of Section 2, that *every* rational (actual) athlete will (should) prefer a ban to complete permission. Instead, I need only that *most* rational (actual) athletes will (should) prefer a ban to complete permission. The argument based on satisfying these preferences still goes through.

That was the easiest case, but perhaps this is too literal a sense of ‘winner’ versus ‘loser’. For example, some baseball players may prefer to dope just to ensure that they make the major league roster. When minimum wage is \$414,000 (and minor league wages are much, much lower), just making the team can count as winning.¹³ In the extreme, what if every professional athlete prefers permissive doping to abstention because there is some (distinct) prize that each athlete wins and that is worth the cost of doping? For some, doping may mean the difference between the home run crown and no crown; for others it may mean the difference between a spot in the starting line up and being on the bench, and yet for others it may mean the difference between a spot on the roster and obscurity in the minor leagues.

I will give two responses to this objection, each based on a distinct symmetry. The first response is suggestive; the second is fatal. First, I grant that if there are many prizes then the chance that dope will help some particular person win at least one of these prizes (that he otherwise would have lost) increases. But this claim has a symmetric counter-

part: if there are many prizes, then the chance that dope will help that same particular person *lose* at least one of these prizes (that he otherwise would have won) also increases. This suggests that we should be sceptical of the idea that mere inflation of prizes will justify permissive doping.

Now the fatal symmetry: for each person who would win some prize only if doping were permitted, there is some other person who would win it if doping were banned. (There may be some overlap: one and the same athlete may both win prize X only if doping were permitted and also win a distinct prize Y if it were banned.) Thus, even if many athletes prefer permissive doping, because of the broad array of available prizes that doping may help them attain, still for every person who would win a prize if doping were permitted, someone else would win it if doping were banned. In this situation we should favour the anti-dopers. This is not because their preferences are more important, nor because there are at least as many of them, but rather because if we satisfy the anti-doping preference then at least no one will be harmed (physiologically) in the process. (Of course, no matter which policy we adopt, many athletes are harmed in the sense of losing a prize they otherwise would have won.) That is, the asymmetric harm of doping serves as a tie-breaker in the symmetric situation where some would win only if doping were permitted and others would win if not.

The last symmetry claim does not entail that at most half of all athletes will prefer permissive doping to a ban. That does not follow from what I said earlier, that for each person who would win some prize only if doping were permitted, someone else would win it if doping were banned. Imagine a sport with two prizes (e.g. home run crown and batting title) and three athletes, A, B, and C. Suppose C would win both prizes if doping were banned, but if it were permitted then A would win one prize, B the other. We should still ban doping in this case even if all this is common knowledge. In effect, C's desire for the first prize is weighted equally with A's desire for the same prize, C's desire for the second prize is weighted equally with B's desire for that second prize, and the harmful effects of doping serve as a tie-breaker. To get to this conclusion we need not aggregate and then maximize preference satisfaction, as a utilitarian might. Rather, we are choosing how to shuffle a fixed set of prizes between different possible winners. If two shuffling methods are otherwise equally impartial (either everyone is permitted to dope, or no one is), but only one inflicts physiological harm to some of the athletes, we should prefer the other method.

Before moving on, let me again emphasize the following point. The hypothesis that athletes can have special epistemic access to the differential degree of enhancement they get from doping is very implausible. Thus, a sufficient response to the possibility that dope enhances some more than others is that, even if so, no rational (actual) athlete will (should) have good reason to think that he is lucky enough to be especially sensitive to dope. In that case, still every rational (actual) athlete will (should) prefer a total ban to permissive doping.

4. Proves Too Much?

In this section I consider whether the game-theoretic argument for a doping ban might also prove, counter-intuitively, that we should ban other methods of performance enhancement besides drugs.

4.1. Low Stakes

An important qualification of the game-theoretic argument is that it applies most forcefully when the stakes are high. This qualification can be interpreted as a relevant constraint in premise 5 of the formal argument of Section 2: athletic associations have greater fiduciary responsibility towards their athletes when stakes are high than when they are low. Hockey helmets, for example, dramatically increase safety while possibly slightly decreasing competitive performance. In such a situation, fiercely competitive players will opt to go helmetless, which is tragic because there is no competitive advantage to anyone if all players forgo helmets.¹⁴ This is a classic positional treadmill, and it suffices to motivate a requirement that all players wear helmets, a policy the National Hockey League (NHL) instituted in 1979. (Of course, helmet mandates create a moral hazard problem in a way that doping mandates do not: NHL players are predictably much more aggressive and reckless now than they were in the 1970s.¹⁵)

However, positional treadmills are ubiquitous in sports, and not all of them are as extreme as the example of hockey helmets. For example, consider the move to carbon fibre parts in cycling. Carbon fibre is slightly lighter than other materials that might do the same job, such as aluminium, but it is also much more expensive. Should cycling organizations ban the use of carbon fibre? The answer depends on the relative costs involved in implementing and enforcing the ban versus the harms incurred without the ban. Yes, carbon fibre is expensive, but the harm incurred in the positional treadmill of upgrading one's bike to carbon fibre materials is trivial compared to the harm incurred on an athlete's health by taking (large doses of) performance-enhancing drugs. The general point, then, is that the game-theoretic analysis presented in Section 2 is more forceful when the degree of harm involved is large.

I do not mean to suggest that the question of whether to allow carbon fibre reduces to a cost-benefit calculation, for example that a ban on carbon fibre is justified if and only if the costs of the ban are less than the combined costs of upgrading from aluminium to carbon fibre for all the racers. In fact, neither direction of that biconditional is plausible. The point is rather that athletic organizations should always try to eliminate positional treadmills when they can, but especially so when the harms associated with those treadmills are severe.

4.2. Competing Desires

Another important qualification of the game-theoretic argument is that it applies most forcefully when there are no competing considerations; that is the point of the constraint in premise 5 of the argument from Section 2. One especially important consideration that competes against what athletes want is what spectators want. For example, cyclists in the Tour de France may choose to descend down a mountain path more or less aggressively. A faster descent increases the probability of injury but may also increase the chance of winning.¹⁶ When choosing the route of the Tour, its organizing body (the Amaury Sports Organisation) must carefully balance safety to its riders against spectator interest, for example by forgoing overly precipitous downhill descents.

Similarly, spectators would lose interest in sports generally if athletic training were banned, even though training clearly represents an arms race scenario: competitive athletes continually escalate their training regimens at the cost of self-harm (if only because they prefer to sleep in and watch television), but with minimal relative payoff

because their competitors are also training more as well. However, as with safe versus unsafe doses of dope, we can also ask whether training should be limited, if not outright banned, on game-theoretic grounds. Perhaps an intermediate position is desirable, where spectators still retain interest even though training is restricted though not banned outright?

No, unrestricted training might still be justified, even though it is a positional treadmill that leads to enormous self-sacrifice. Note that at least the first two points I made in Section 3.1 about safe versus unsafe doses of drugs apply to training as well. First, there is no sharp cut-off between harmful and non-harmful levels of training. Rather, the degree of harm increases monotonically with the amount of training. Second, even if there is a sharp cut-off between harmful and non-harmful levels of training, the position of that cut-off is both subjective (different for different people and in different situations) and controversial (open to normative dispute even when all the descriptive facts are known).

The third point I made about safe versus unsafe dosages of dope—that the marginal cost of moving from a cap on drug doses to a complete ban is negligible—is less clearly applicable here. That is because the way we test for drug use differs from the way we restrict training: for drugs we acquire urine or blood and measure chemical concentrations, and it is no more expensive to look for zero in the results than for some non-zero number. For training we have to do some other thing (hire personal chaperones?) which may become more expensive as we restrict the permitted amount of training ever more tightly.

In any case, there is another important difference between training and dope. We clearly want to *limit* the use of performance-enhancing drugs at least somewhat. Once we acknowledge that some limit on dope is desirable, the important question becomes how far to limit it, and in that case the negligible marginal cost of a total ban is relevant. In contrast, we clearly want to *allow* some amount of training. First, athletes enjoy some aspects of training, for example playing catch or hitting golf balls on a beautiful day, so to preclude them from engaging in any training whatsoever seems perverse. Second, think of how fans would be turned off from baseball if every player just relied on raw ability and never engaged in any training, practice, or conditioning. Indeed, that most spectators prefer to allow at least some training may also be an indirect reason for athletes to be in favour of that as well, as spectators indirectly pay their salary. (Note also that this appeal to spectator preferences is not nearly as controversial as the assertion that spectators would be turned off to sport if athletic doping were ubiquitous. My claim is rather, and much more plausibly, that they would be turned off to sport if training were banned.)

Thus, given that we want to allow at least some amount of training, our initial marginal cost question is whether to restrict training at all, and here the marginal cost is of course very large. Thus, (a) we incur a significant start-up cost in agreeing to restrict training at all, (b) the marginal cost of ever tighter restrictions need not fall (as they would for drug tests), and (c) as we increase such restrictions spectator backlash will increase as well. In such a situation, staying with no training restrictions at all may well be appropriate.

Thus, we must complicate our justification for a total doping ban again. I am trying to explain why we should ban doping. But the answer to the question ‘why should we ban doping?’ depends on the contrast class. Are we asking ‘why should we *ban* doping

rather than *allow* it?', or are we asking, for example, 'why should we ban *doping* rather than ban *training*?'. The answer to the first question—why we should ban doping rather than allow it—is the game-theoretic argument I gave in Section 2. The answer to the second question—why we should ban doping rather than training—is that banning training would incur considerable costs, including a significant spectator backlash; not so if we ban doping.

Again, the argumentative strategy I use here is instructive not only where it succeeds but also where it fails. Spectator preferences and other costs can be outweighed if the risks to athletes are sufficiently extreme. For example, the NFL and its players' union have recently discussed whether it would be appropriate to restrict practices in order to reduce the risk of concussion.¹⁷ And in collegiate sports, an opportunity cost of additional training is less education, which is quite significant for student-athletes. The NCAA has some fiduciary responsibility to its student-athletes, so it predictably imposes strict limits on practices and games.¹⁸

5. Taxonomy

In this section I discuss two related objections, disguised as questions of taxonomy. I argue that a game-theoretic doping ban is neither *paternalistic* nor *coercive*, at least in the objectionable senses of those terms.

5.1. Paternalism

The game-theoretic rationale for a doping ban is more vulnerable to the charge of paternalism than is a ban on grounds of, say, unnaturalness. That is because unnatural things are not necessarily harmful, whereas an essential premise of the game-theoretic argument is that dope is harmful. Still, the game-theoretic argument can easily evade the following form of argument: 'A doping ban is paternalistic; paternalism is wrong; therefore we should not ban doping.' It evades that line of reasoning because the reasoning trades on two different senses of 'paternalistic'. One sense applies to the game-theoretic argument but is not wrong; the other would be wrong but does not apply to the argument.

The game-theoretic argument is paternalistic in the sense that it is motivated out of concern for the welfare of the athletes involved. However, not all instances of acting to promote the welfare of someone other than yourself are thereby objectionable. In this sense of 'paternalistic', the game-theoretic rationale for a doping ban is paternalistic but not wrongfully so.

Now, most people think a paternalistic act must in some sense be contrary to the victim's will.¹⁹ And if a doping ban went against the will of the athletes, that might be wrongful. But, of course, athletes prefer a total ban to total permission, so this further constraint on the definition of paternalism is unhelpful to my critics.

Still, even if most players want a ban, such a ban might seem inappropriate because it is being implemented *coercively*: players who disobey the mandate are punished, and they have no choice over the matter. And if we coercively prevent athletes from doping when they want to, that might be paternalistic in an objectionable sense.²⁰ However, a doping ban is not coercive in any objectionable sense, as I will now show.

5.2. *Coercion*

The best way to think about coercion is in terms of whether a given proposal is a threat or an offer.²¹ A proposal is a *threat* when it forecloses opportunities, relative to some larger comparison class of opportunities. For example, the proposal ‘your money or your life’—made by someone with a loaded gun—forecloses the opportunity of continuing along your way with *both* your money and your life. That is a threat, and therefore coercive. In contrast, a proposal is an *offer* when it expands your opportunity set to include new possibilities. For example, the proposal ‘please come work for me’, made by someone other than my current boss, gives me another opportunity I previously lacked. That is an offer, and therefore not coercive.

The key question for us, then, is whether a doping ban shrinks versus expands an athlete’s opportunity set, and this depends on the relevant comparison class of opportunities. This comparison class may be temporal (what opportunities you had initially), counterfactual (what opportunities you would have had in some non-actual situation), or moral (what opportunities you should have had). Which of these types of comparison class we should favour, if any, is a question of some theoretical interest.²² In what follows, however, I will suppress reference to this taxonomy and simply consult intuitions about appropriate comparisons.

On one way of thinking, a ban eliminates one of the two antecedent possibilities, ‘dope’ and ‘abstain’. That shrinks the athlete’s opportunity set, so a proposal of the form ‘stop doping or else suffer some sanction’ is, on this interpretation, coercive. Worse, if every prudentially rational athlete prefers to dope when given the chance, the ban is especially egregious because it precludes the possibility that rational athletes most favour.

However, that way of thinking is incorrect. A more accurate way of describing a doping ban is as terms of employment, in other words as folded into an antecedent offer, such as ‘please join this league where you will not be allowed to dope’. On this way of thinking, a doping ban is no more wrongful than is requiring players to wear certain colour socks for home games—the ban on dope and on blue socks are both just conditions to which the athlete agrees when he chooses to play for the relevant league. Thus, these terms of employment are still parts of a larger offer rather than being independent, narrow threats; the larger offer is just perhaps less attractive than it might otherwise have been.²³

Moreover, we must remember two points. First, a doping ban is implemented for all players. Thus, the relevant proposal for any given athlete is not that he alone forgo dope. It is that *everyone* forgoes dope, including his competitors. In effect, the ban creates a new option for the athlete that would not otherwise exist, the option that no one dopes. Yes, the ban also mandates that option, but we must also remember this second point: athletes prefer a universal doping ban to universally permissive doping. So, in fact, this new option—that no one dopes—is the most preferred of all the (impartial) options.

Notice that this last point—that a complete ban on doping is more attractive to athletes than is complete doping permission—is not necessarily available to other doping ban arguments. For example, a naturalness argument will insist that unnatural things should be banned even if they benefit all athletes independently of their athletic performance. Consider a hypothetical wonder drug that increases strength and lengthens life

with no harmful side effects. Clearly, all athletes would be in favour of permission to take this wonder drug, yet a naturalness argument would still recommend a ban on its use, and that might seem wrongfully coercive. The game-theoretic argument, in contrast, would not license that ban.

6. Limits

In this final section, I will mention some of the limits of my argument. First, as just discussed in Section 5.2 (the hypothetical wonder drug that increases strength and lengthens life, with no harmful side effects), the game-theoretic argument does not justify banning non-harmful drugs. The desire to forbid athletes, or indeed anyone, from using a substance which only helps them and does not harm them in any way (including the cost of opportunities forgone to acquire the substance) strikes me as perverse. Thus, I view this specificity of the game-theoretic argument as one of its virtues. At the very least, if we want to justify a ban on non-harmful (beneficial) substance it will have to be on the basis of some argument other than the one I have presented here.

Second, and as also already mentioned, the game-theoretic argument does not license a ban if reliable testing or enforcement is unfeasible. Unfeasibility can be due to technological limits in detection, but it may also stem from financial limits. That is why I restricted my thesis to *large* athletic organizations, though the more accurate and important criterion is wealth. Indeed, all the athletic organizations I have mentioned in this paper have deep pockets, and my argument does not apply to significantly poorer organizations, for example a local softball league.

Here is a third possible restriction. Suppose testing for doping agent X is feasible, yet X has a cheap and readily available substitute, Y, for which testing is unfeasible. For example, X might be an anabolic steroid and Y its precursor, or X might be EPO doping where Y is blood doping (storing one's own blood for transfusion right before a race). In such a situation, the inability to test for Y may mean that a ban on X is futile, as athletes will just switch from using X to using Y. In such cases, perhaps a ban on X is, all things considered, unjustified.

Finally, I would like to emphasize again a few theses in the neighbourhood of mine—we should ban athletic doping—that the game-theoretic argument does not support. I argued that athletic organizations *should* do something—namely, ban doping in their sports—but of course that is compatible with the thesis that they have the right not to do that thing. In other words, while it is not immorally coercive for a private athletic organization to ban doping within its sport, it might be immorally coercive for the state to mandate that (some) private athletic organizations ban doping within their sports. I take no stand on this controversial issue.

Similarly, to say that private athletic organizations should ban doping within their sports is not to say that athletes should not dope. Indeed, my argument presupposes that it is often instrumentally rational for athletes to dope, when given the chance.

Along similar lines, the game-theoretic argument pertains only to competitive sports. It does not pertain to, say, mountaineering for pleasure. Yes, high altitude hiking is (more) dangerous (than low altitude hiking), but I have no objection to such hiking on its own. What would be objectionable—more like 'tragic', in the game-theoretic sense—would be if two people intent on beating each other in a race both went to high

altitudes to train, thereby both incurring risk of harm without any gain relative to the other. That is when the game-theoretic argument becomes relevant.

Eric Chwang, Department of Philosophy, University of Colorado at Boulder, Boulder, CO 80309-0232, USA. Eric.Chwang@Colorado.edu

Acknowledgments

Thanks to Jon Dang and Paul Bowman for helping talk through the issues with me. Thanks to an audience at the University of Colorado at Boulder's Center for Values and Social Policy for their helpful feedback on a presentation of this material. Special thanks to David Boonin, Chris Heathwood, and Brian Talbot for written feedback on earlier drafts. And Michael Huemer, whose amazingly brilliant comments made this paper into the fine specimen of philosophical insight that it is today. (Mike told me to insert that last sentence fragment.)

NOTES

- 1 For a sceptical view of testing feasibility, see B. Kayser & A. Smith, 'Globalisation of anti-doping: The reverse side of the medal', *British Medical Journal* 337 (2008): 85–87.
- 2 R. L. Simon, *Fair Play: The Ethics of Sport*, 3rd edn. (Boulder, CO: Westview Press, 2010), pp. 85–87 comes close to articulating the argument, though not in game-theoretic terms, as I do, and he does not respond to the same objections I do. His presentation also differs in other important respects from mine. For example, Simon asks what impartial athletes would choose under a veil of ignorance, whereas my argument is compatible with the possibility that athletes act in self-interest and that they have complete information about the differential effects of dope on them and their competitors.
- 3 See, for example, M. Sandel, *The Case Against Perfection: Ethics in the Age of Genetic Engineering*, (Cambridge, MA: Belknap Press, 2007); N. Dixon, 'Performance-enhancing drugs, paternalism, meritocracy, and harm to sport', *Journal of Social Philosophy* 39 (2008): 246–268; N. Fost, 'Steroid hysteria: Unpacking the claims', *Ethics Journal of the American Medical Association* 7 (2005): <<http://virtualmentor.ama-assn.org/2005/11/pdf/oped2-0511.pdf>> (last visited 22OCT11); Kayser & Smith op. cit.; and Brent M. Kioussis, 'Philosophy on steroids: Why the anti-doping position could use a little enhancement', *Theoretical Medicine and Bioethics* 29 (2008): 213–234.
- 4 See R. H. Frank, *Choosing the Right Pond: Human Behavior and the Quest for Status*, (New York: Oxford University Press, 1985), chapter 7.
- 5 See B. Tindall, 'Better Playing Through Chemistry', *The New York Times* (17 October 2004), available at: <<http://www.nytimes.com/2004/10/17/arts/music/17tind.html>> (last visited 22OCT11).
- 6 Sandel op. cit., for example, shifts between all three forms of argument.
- 7 See, for example, Dixon op. cit.
- 8 Ibid.
- 9 For some suggestive data on the role model effect on adolescents, see J. R. Hoffman A. D. Faigenbaum, N. A. Ratamess, R. Ross, J. Kang & G. Tenenbaum, 'Nutritional supplementation and anabolic steroid use in adolescents', *Medicine & Science in Sports & Exercise* 40 (2008): 15–24.
- 10 See, for example, C. L. Carr, 'Fairness and performance enhancement in sport', *Journal of the Philosophy of Sport* 35 (2008): 193–207. Many anti-doping articles contain a mixture or combination of some or all the above sorts of arguments, for example C. Lenk, 'Is enhancement in sport really unfair? Arguments on the concept of competition and equality of opportunities', *Sports, Ethics and Philosophy* 1 (2007): 218–228.
- 11 See M. S. Juhn, 'Popular sports supplements and ergogenic aids', *Sports Medicine* 33 (2003): 921–939.
- 12 See Juhn op. cit., p. 935 and S. Achar, A. Rostamian, & S. M. Narayan, 'Cardiac and metabolic effects of anabolic-androgenic steroid abuse on lipids, blood pressure, left ventricular dimensions, and rhythm', *The American Journal of Cardiology* 106 (2010): 893–901.

- 13 The figure is current for 2011. See <<http://mlb.mlb.com/pa/info/faq.jsp#minimum>> (accessed 22 October 2011).
- 14 See Frank op. cit., chapter 7.
- 15 See A. Schwarz, 'A case against helmets in lacrosse', *The New York Times* 16 February (2011), available at: <<http://www.nytimes.com/2011/02/17/sports/17lacrosse.html>> (accessed 22 October 2011).
- 16 The example is from Kiouss op. cit., p. 216.
- 17 See M. Maske, 'NFL, union discuss practice restrictions to prevent concussions', *The Washington Post* 2 June (2010), available at <<http://views.washingtonpost.com/theleague/nflnewsfeed/2010/06/nfl-union-committees-discuss-possible-practice-restrictions-to-prevent-concussions.html>> (accessed 22 October 2011).
- 18 See, for example, the first sentence of Article 17 (Playing and Practice Seasons) of the 2011–2012 NCAA Division 1 Manual (July 2011), available at <<http://www.ncaapublications.com/productdownloads/D112.pdf>> (accessed 22 October 2011): 'A member institution shall limit its organized practice activities, the length of its playing seasons and the number of its regular-season contests and/or dates of competition in all sports, as well as the extent of its participation in noncollegiate-sponsored athletics activities, to minimize interference with the academic programs of its student-athletes.'
- 19 For example, the first sentence of G. Dworkin, 'Paternalism', *The Stanford Encyclopedia of Philosophy* (Summer 2010 Edition), Edward N. Zalta (ed.): <<http://plato.stanford.edu/archives/sum2010/entries/paternalism/>> (accessed 22 October 2011), with my emphasis: 'Paternalism is the interference of a state or an individual with another person, *against their will*, and defended or motivated by the claim that the person interfered with will be better off or protected from harm.'
- 20 For more on paternalism with and without coercion, see B. Gert and C. M. Culver, 'Paternalistic behavior', *Philosophy and Public Affairs* 6 (1976): 45–57.
- 21 The classic exposition of this is R. Nozick, 'Coercion', in S. Morgenbesser, P. Suppes & M. White (eds) *Philosophy, Science and Method: Essays in Honor of Ernest Nagel* (New York: St. Martin's Press, 1969), pp. 440–472. Also see A. Wertheimer, *Coercion* (Princeton, NJ: Princeton University Press, 1987).
- 22 For example, both Nozick op. cit. and Wertheimer op. cit. discuss this question at some length.
- 23 See H. G. Frankfurt, 'Coercion and moral responsibility', in T. Honderich (ed.) *Essays on Freedom of Action*, (London: Routledge & Kegan Paul, 1973) for more on this distinction between threats and less attractive offers.