MODELLING DRUG TAKING IN SPORT

CLAIRE BAUDOUIN

Magdalen College, University of Oxford
Michaelmas Term, 2015

Submitted for the degree of Doctor of Philosophy in Economics

ABSTRACT

This thesis considers optimal policies in response to doping in sport. Doping is an area of increasing public concern as growing numbers of athletes, and the organisations responsible for testing them, are implicated in doping scandals. These scandals affect not only those interested in sport, but also society as a whole since athletes are role models. Consequently it is important to consider optimal responses to doping.

The second chapter considers the optimal ban for a welfare maximising anti-doping agency (ADA) to impose on an athlete who tests positive for performance enhancing drugs. Previous authors have assumed that maximal deterrence is optimal, but this chapter shows that when the form of the punishment is a ban from competition, maximal deterrence may no longer be optimal. The optimal punishment is found to depend on factors such as the prevalence of doping and the variance of earnings.

The third chapter examines situations in which ADAs may not wish to uncover doping. In this case, if consumers cannot observe the frequency with which athletes are tested, a no-cheating equilibrium does not exist. I find two mechanisms which can lead to a no-cheating equilibrium: consumers who can observe the frequency of testing and an external testing agency, such as WADA, which can retest athletes and punish the ADA if it failed to detect doping.

The final chapter examines whether increasing the frequency of testing deters athletes from doping. Since data is not available to analyse this problem directly, I use the relationship between testing and Olympic performance to infer the relationship between testing and doping. The results suggest that while in some sports, such as athletics, carrying out more tests does deter athletes from taking drugs, in others, such as cycling, there is no evidence of a negative relationship between testing and doping.

Word Count: 45,000

ACKNOWLEDGMENTS

I would like to thank my supervisors Tom Norman and Peyton Young for their help and guidance throughout the research and writing of this thesis. Jim Malcomson also provided me with much advice as my MPhil supervisor. I am grateful to Stefan Szymanski who acted as my supervisor during my visit to the University of Michigan where I began work on my empirical chapter. I would also like to thank Ian Crawford who was very generous with his time and gave me valuable feedback on my empirical chapter.

I am grateful to the Economic and Social Research Council (ESRC) for their financial support. My doctoral studies were funded by the ESRC grant ES/J500112/1.

Finally I would like to thank my family and all those who have made my time in Oxford so enjoyable.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 INTRODUCTION

1.1 OVERVIEW

Doping in sport has become an increasing concern in recent decades with many high profile athletes being implicated in doping scandals. Doping not only undermines the enjoyment of sport for those who watch it, but also has a wider social impact since many athletes are role models. Given this, and the key role which incentives play in the problem of cheating, it is not surprising that there is a growing literature on the economics of doping.

One of the most striking features of the current anti-doping system is its complexity. The World Anti-Doping Agency (WADA) is the chief authority in anti-doping. It is responsible for setting the rules relating to doping including the punishments for doping offences and the list of banned substances. However, WADA itself carries out very few tests on athletes. In 2014 it carried out less than 0.3% of total tests (WADA, 2015*a*). Instead, these tests are carried out by subsidiary organisations. There are two main type of subsidiary testing organisations: national anti-doping organisations (NADOs) and international federations (IFs) which are sport-specific organisations. For example, there are two main organisations which can test a British cyclist: UK Anti-Doping and the International Cycling Union. Subsidiary organisations must accept the rules imposed by WADA but are free to choose the frequency with which athletes are tested.

Chapters 2 and 3 of this thesis use theory to examine different aspects of the anti-doping system. The second chapter considers the social welfare maximising ban for WADA to impose on an athlete who tests positive. Particular emphasis is placed on how the characteristics of sports, such as the variance of earnings, determine the optimal ban. However, even if WADA chooses the optimal ban, the organisations responsible for conducting testing may lack the incentives to uncover cheating. National and sport-specific testing agencies may have little interest in discovering cheating among their athletes since this is likely to be damaging to their reputation. Therefore the third chapter considers how the incentives of non-benevolent testing agencies can be changed so that athletes are tested sufficiently frequently that they are deterred from doping.

The fourth chapter uses empirics to examine whether increased testing deters athletes from taking drugs. If athletes believe they are at risk of testing positive when they are tested, we would expect increased testing to deter doping. If testing does deter doping then this suggests that a relatively simply solution to the problem of doping in sport is to increase testing. On the other hand, if it does not, we may wish to consider redesigning the system.

1.2 RELATED LITERATURE

There is a growing body of literature relating to the economics of doping.[1] This literature is predominantly theoretical and can be divided into two strands. The first focuses on the interaction between athletes and is mainly concerned with when doping and non-doping equilibria are most likely to arise. The second considers the relationship between an athlete and an anti-doping agency (ADA). This strand focuses on how the presence of ADAs may help to deter athletes from doping.

With regard to the first strand, Breivik (1992) was the first to consider doping within a game theoretic framework. He focuses on symmetric, two-athlete games and shows that when athletes view winning as the most important goal, the result is a prisoner's dilemma situation in which both athletes dope even though the no-doping equilibrium is a Pareto improvement. Haugen (2004) also models the two-athlete game as a prisoner's dilemma while Haugen et al. (2013) extend the analysis to the n-athlete game. Bird and Wagner (1997) study doping as a common pool resource dilemma and suggest that a potential solution is to have no banned substances but instead to insist athletes record all drugs

1 A good overview of this literature is provided in Dilger, Frick and Tolsdorf (2007)

they are using in a drug diary. The authors suggest that social norms may then develop which would lead athletes to choose not to dope. In a similar manner, Eber (2008) extends Haugen (2004) by introducing fair-play norms. He finds a no-doping equilibrium then exists, although this equilibrium is not unique.

Berentsen and Lengwiler (2004) and Kräkel (2007) both allow for heterogeneous athletes. Kräkel allows athletes to choose not only doping, but also a legal input such as training. He shows that a no-doping equilibrium is more likely to occur if the degree of heterogeneity increases. Meanwhile, Berentsen and Lengwiler (2004) consider an evolutionary game. They find that under some circumstances, there may be cycles of doping and clean sport. Additionally, for some parameter values highly talented athletes may be more likely to cheat than less talented ones.

Berentsen (2002) can be seen as being at the intersection of the two strands of literature. Berentsen first analyses a two player game where one player is more talented than the other. He then introduces sanctions for testing positive, firstly when the sanction cannot be dependent on ranking (the *status quo*) and secondly when it can be. Berentsen finds that a no-doping equilibrium cannot be implemented under all circumstances when the sanction is not rank dependent. However, if the winner is fined a specific amount while the loser is not punished if he tests positive, a no-cheating equilibrium always exists.

Berentsen (2002) does not directly analyse the problem from the perspective of an anti-doping agency though. Kirstein (2012), on the other hand, falls firmly into the second strand of the literature since he considers an inspection-style game where the testing technology is imperfect in the sense that false positives and false negatives can arise. He assumes that no incentive problems arise for the tester and the tester is benevolent. Kirstein finds that if testing provides an imperfect signal, an equilibrium in which the athlete cheats with zero probability cannot be achieved. Berentsen et al. (2008) also consider an inspection game with a benevolent testing agency. However, they do not consider imperfect tests but instead introduce a whistle-blowing stage whereby the loser can blow the whistle on the winner. The authors find that the addition of this stage reduces the frequency of cheating.

Maennig (2002) does not explicitly analyse the interaction between an athlete and ADA, but instead draws on Becker (1968) to analyse the decisions made by corrupt athletes or officials. Eber (2002) and Emrich and Pierdzioch (2013*a*) both also focus on corrupt officials, but more specifically on the incentives for national anti-doping agencies to tolerate a higher level of doping than would be the case if anti-doping policies were co-ordinated across countries. However, neither of these papers explicitly analyse the choice problem of athletes.

Finally, Buechel et al. (2014) extend the second strand of the literature by also introducing consumers as an additional player in the game between an athlete and anti-doping agency. They assume a benevolent ADA but show that if consumers leave after a doping scandal, even a benevolent ADA will not carry out any tests and the only equilibrium outcome is that the athlete dopes.

There are far fewer empirical papers concerning doping in sport. This is unsurprising given the illegal nature of doping. Those empirical studies which do exist tend to attempt to estimate the prevalence of doping among elite athletes. While the results of tests carried out by ADAs generally indicate that only 1-2% of athletes are doping, empirical studies find a significantly higher proportion are doing so (e.g. Werner Pitsch et al., 2007; Pitsch and Emrich, 2012; Striegel et al., 2010). Meanwhile, Dilger and Tolsdorf (2005) analyse the characteristics of athletes who have tested positive. They find that athletes who have been caught cheating participated in significantly fewer races than comparable clean athletes. Finally, there is some empirical evidence regarding how consumers react to doping scandals. Van Reeth (2012) shows that TV viewership of the Tour de France declines after a new case of doping is revealed, while Cisyk and Courty (2015) find that fan attendance at Major League Baseball matches declines after a doping violation is announced.

1.3 SUMMARY OF CHAPTERS

Chapters 2 and 3 of this thesis are predominantly theoretical and are most closely aligned with the strand of existing literature which considers the relationship between an athlete and an anti-doping agency. Chapter 2 draws on the approach developed by Becker (1968) to consider the optimal punishment for a benevolent organisation, such as WADA, to impose on an athlete who tests positive. The existing literature on doping has generally not considered this issue, since it is usually assumed that no-doping is the desirable outcome and the optimal punishment is then that which supports the no-doping equilibrium. However, this analysis fails to take into account a key aspect of punishments for athletes. The existing literature models the punishment as a lump-sum fine. In reality though, athletes who test positive are banned for a proportion of their career. This is important for two reasons. Firstly, when an athlete is banned from competing, society cannot benefit from the athlete's performance and this reduces social welfare. Secondly, this structure of punishment means that athletes with higher career earnings have a larger absolute cost of testing positive than athletes with lower earnings.

When the punishment is a ban from competition, a no-doping equilibrium may not be optimal. Instead, in choosing the optimal punishment, WADA must trade-off two counteracting forces. Firstly, an increase in the ban is beneficial for social welfare because it reduces the proportion of athletes taking drugs and society suffers a social loss if a role model is found to have been cheating. Secondly though, increasing the ban is damaging to social welfare since it reduces the expected career length of athletes.

Results are derived using comparative statics and simulations. The optimal punishment is found to depend on the characteristics of the sport in question. This chapter therefore adds weight to arguments against WADA's current policy of having standardised sanctions across sports. The results suggest that sports with a greater variation of earnings should have a harsher ban. Scully (1995) has found that team sports are likely to have lower variance of earnings than individual sports. This suggests that *ceteris paribus* sports such a hockey should have a shorter ban than sports such as athletics. The results also suggest that as doping becomes more endemic in a sport, at first it is best to increase the punishment in an attempt to deter athletes from doping but eventually it becomes optimal not to attempt to deter doping.

The second chapter focuses on the role of WADA in the anti-doping system and therefore abstracts from the testing process in order to concentrate on optimal punishments. However, if the agencies responsible for testing are not benevolent, even if WADA sets the punishment optimally, social welfare will not be maximised. Therefore the third chapter explores the possibility that the agencies responsible for testing athletes may not want to catch cheats. For example, Lance Armstrong has claimed that the International Cycling Union helped cover up evidence that he tested positive at the 1999 Tour de France (Huffington Post, 2013). In general agencies may not wish to see their reputation tarnished by uncovering doping by athletes associated with their country or sport. This chapter therefore examines ways in which non-benevolent anti-doping agencies can be incentivised to test with a sufficiently high frequency that doping is deterred. The chapter draws on Buechel et al. (2014), but examines a non-benevolent as opposed to a benevolent ADA and also models consumers as fully rational players.

Two main mechanisms are studied for incentivising ADAs. The first is if consumers are players in the game, engage in Bayesian updating and are only prepared to participate if they are sufficiently confident that the athlete did not cheat. This can result in a no-doping equilibrium, but only if consumers are able to observe the frequency with which athletes are tested. If anti-doping agencies are not prepared to publish this information, an alternative is to have an external agency, such as WADA, which randomly retests athletes and punishes the ADA if doping is detected when the ADA failed to detect this.

One of the key ideas in the theoretical analysis of doping is that increased testing deters athletes from taking drugs. If testing does deter doping then a relatively simple solution to the problem of doping is to test athletes more frequently. However, this will only work if athletes believe there is a risk of testing positive when they are tested. Therefore my fourth chapter uses empirics to look for a relationship between testing and doping. Unfortunately, this issue cannot be examined directly since we rarely know whether or not an athlete chose to take drugs. Instead, this chapter takes a more indirect approach, using the relationship between testing and Olympic performance to infer the relationship between testing and doping.

The results suggest than in some sports, such as athletics and wrestling, increased testing does deter athletes from taking drugs. In other sports where doping is believed to be widespread though, such as cycling, there is no evidence that this is the case. This suggests that in certain sports, drug tests may either be incapable of detecting the substances which athletes are using, or as is suggested in my third chapter, in the current system ADAs may lack incentives to uncover cheating.

This thesis suggests implications for anti-doping policy. The second chapter shows that since the punishment structure in sport is a ban from competition, the optimal anti-doping policy may not be maximal deterrence. Instead, WADA should consider setting different bans for different sports, in order to take account of their differing characteristics. The third chapter suggests that increased transparency is necessary to ensure that non-benevolent anti-doping agencies do not attempt to cover up doping scandals. This analysis is particularly pertinent given the recent report into corruption in Russian athletics (Pound et al., 2015). Finally, the fourth chapter suggests that while in some sports increased testing may be sufficient to deter doping, in other sports, the problem appears to have deeper roots.

2 OPTIMAL BANS FOR DOPING OFFENCES

2.1 INTRODUCTION

Floyd Landis and Matt Dumontelle both tested positive for banned substances in 2006 and both received a ban of two years. Landis was an American cyclist who tested positive for performance enhancing drugs after a spectacular stage in the 2006 Tour de France while Dumontelle was a Canadian curler who tested positive after the final of the World Curling Championships. The fact that both athletes received the same ban despite competing in very different sports is in line with the World Anti-Doping Agency's policy that sanctions should be standardised across sports. However, in its World Anti-Doping Code, WADA notes that harmonisation of sanctions across sports is one of the most controversial areas of anti-doping (WADA, 2015*b*). WADA's arguments in favour of harmonisation focus on fairness and the need to limit flexibility in order to prevent some sporting organisations from exploiting such flexibility in order to behave more leniently towards drug takers (WADA (2015*b*), p.78).

This chapter considers how the optimal punishment for each sport should be determined if WADA were to relax its policy of harmonisation. The characteristics of different sports vary widely and as would be expected, I find these characteristics determine the optimal punishment from the perspective of society. This chapter therefore suggests arguments against the continuation of standardisation. Sport-specific sanctions could be implemented without raising concerns about increased flexibility if WADA itself set the sanction for each sport. However, it is not possible to examine issues of fairness within the model presented here. Ultimately policy-makers must decide whether they believe the moral arguments for harmonised sanctions outweigh the potential social benefits from moving to sport-specific sanctions.

In the model presented here, WADA wishes to maximise a social welfare function, where the choice variable is the percentage of an athlete's career for which the athlete is banned from competing if he tests positive for taking drugs. Society values higher quality performances, but can only benefit from an athlete's performance if he has not been banned from competing. The existing literature models the punishment for an athlete who tests positive as a lump sum fine. However, this ignores the fact that when an athlete is punished, society is also harmed since it does not benefit from the athlete's performance for the duration of his ban. It is this aspect of the punishment which means that in the model presented here, society may not wish to maximally deter doping.

In order to model the interaction between WADA and the athlete, attention is restricted to a single athlete. Consequently, this model is most suited to sports where absolute performance plays an important role. Absolute performance is likely to be important in sports where athletes compete against a number of competitors simultaneously rather than head to head. For example, in sports such as athletics, athletes are likely to be at least partially paid on the basis of the times they achieve and society derives greater pleasure from watching the athlete when the athlete achieves better times. If athletes are paid purely on the basis of their absolute performance, it is plausible to focus on only a single athlete for two reasons. Firstly, the decision of other athletes as to whether or not to take drugs does not impact on the advantages and disadvantages of drug taking for an individual athlete and it is not necessary to model interdependence between the actions of athletes. Secondly, changes in the performance of this individual athlete do not impact on the benefit which society obtains from other athletes and therefore attention can be focused on the interaction between WADA and a single athlete. It is clearly a simplification not to model interactions between athletes in the decision by the athlete of whether or not to take drugs. However, the decision by athletes in this model can be seen as a reduced form of such an interaction and others, such as Maennig (2002), have also abstracted from relative performances.

All athletes have the same baseline ability if they do not cheat although performance also has a stochastic element. Athletes differ in their gain from drug-taking though. This reflects the fact that different athletes have different physiological responses to the same drugs. Athletes know their own type but WADA only knows the underlying distribution from which types are drawn. Athletes decide whether or not to take drugs based on the probability of testing positive if they do or do not take drugs and the sanction set by WADA. As the sanction increases, fewer types of athlete will choose to take drugs.

In choosing the optimal punishment WADA must trade-off two conflicting forces. Firstly, if the length of ban increases while the proportion of athletes taking drugs is held constant, there is a social loss since if an athlete tests positive, he is banned for a longer proportion of his career and society does not benefit from the athlete's performance when he is banned. This is the punishment channel. Secondly though, if we allow the proportion of athletes doping to decline as the ban increases, but hold constant the period for which society does not benefit from performance if the athlete tests positive, social welfare improves. This is because if the athlete tests positive, society suffers a lump sum cost reflecting the damage when it emerges a role model took drugs. When fewer athletes dope, the expected lump sum loss due to an athlete testing positive declines. This is the proportion channel. At the optimum, these two effects exactly offset one another.

Results are derived using both comparative statics and simulations. The comparative statics show that when the social cost of drug taking increases, the optimal punishment should increase. This suggests that more popular sports should have a longer ban for doping. However, the comparative statics are ambiguous about the effect of a change in the mean or variance of the gain from doping. Consequently, simulations are carried out using plausible parameter values. The results suggest that sports with a greater variation of earnings should have a harsher punishment. Scully (1995) has found that the variance of earnings differs widely across sports and in particular that team sports generally have a lower variation of earnings than individual sports. This suggests that sports such as curling should have shorter bans than sports such as athletics. The mean gain from doping dictates the proportion of athletes who would be willing not to dope for any given punishment. Therefore, it can be seen as a measure of the extent to which a 'culture of doping' exists in a sport. The results suggest that as doping becomes more endemic in a sport, while at first it is optimal to increase the punishment to counteract this, eventually it is better not to attempt to deter athletes from doping.

2.2 RELATED LITERATURE

Becker (1968) was the first paper to apply economics to the topic of crime and deterrence. In this paper Becker analyses the optimal policies in the sense of those which minimise the social loss from criminal offences. He shows that the optimal punishment is greater the greater the damage of the offence, a conclusion which is echoed in my model. Stigler (1974) builds on Becker and in particular emphasises that complete deterrence from crime is generally neither possible nor optimal since enforcement is costly. Stigler also differs from Becker since he does not believe the gain to the offender should enter into the social welfare function.

This chapter draws on these papers and applies the approach developed in them to the problem of doping in sport. However, applying this style of model to the question of the optimal punishment for doping offences requires some substantial adaptations. Firstly, in the case of doping offences, society can gain from the offence because society values higher quality performances by athletes. This is in contrast to Stigler's view that when an offence is committed, society does not benefit from the offender's gain. Secondly, the form of punishment which is used in sport is a ban from competing for a given number of years. Becker (1968) suggests that when the punishment is a fine, the social cost is approximately zero since the fine can be viewed as a transfer. This is not the case when the punishment is a ban from competition. In this case, a harsher punishment is costly to society not because society must bear the cost, for example, of running a prison, but because the social benefit from watching the athlete compete is foregone when the athlete is banned from competition. Finally, since athletes can always choose an alternative career but society benefits from athletes competing, it is necessary to consider a participation constraint.

In the current sports economics literature, Maennig (2002) is one of the few authors to use a crime approach to analyse drug taking in sport, rather than the more standard game theoretic approach. Maennig analyses the factors which influence whether or not an individual athlete chooses to take drugs, but does not examine the optimisation problem of WADA. Most other papers (e.g. Breivik, 1992; Haugen, 2004) focus on the existence of a no-doping equilibrium and assume that such an equilibrium is desirable. Such papers generally do not consider an anti-doping agency or the optimal sanction for such an agency to impose. Additionally, in the existing literature, the punishment for testing positive for drug taking is always assumed to be a lump sum fine, rather than the more appropriate assumption used throughout this paper that athletes are banned for a proportion of their careers.

2.3 MODELLING

2.3.1 *Set-up*

In this model there is an anti-doping agency which wishes to maximise a social welfare function. The anti-doping agency achieves this by choosing the proportion of an athlete's career, n, for which the athlete is banned from competing if he tests positive for drug taking. Consequently it must be the case that $n \leq 1$. All parameters except n are exogenous for WADA.

It is assumed that an athlete has the same expected earnings in every year of his career and therefore the ban imposed by WADA is equivalent to the athlete losing proportion n of his expected lifetime income from competing. This structure of punishment reflects the idea that the sanction imposed by anti-doping agencies is not a lump-sum fine, but rather a ban from competing for a given number of years. An important feature of this punishment structure which has been ignored in previous analysis is that the punishment is proportional to performance. Since in this model the decision of an individual athlete as to whether or not to take drugs is independent of the decision made by any other athlete, the model can be simplified by assuming that WADA only deals with one athlete.

All athletes have the same baseline ability, a, and produce a performance of quality $a + \epsilon$ if they do not take drugs. a is exogenous and ϵ is mean zero, has support $(-\infty, \infty)$ and is independent of b. ϵ is not realised until after the athlete has made his decision about whether or not to take drugs and reflects the role of chance in an athlete's performance. Athletes of different types derive different benefits from drug taking. If an athlete takes drugs his performance is $ab + \epsilon$. Therefore b is the multiplicative benefit (or loss) in performance that an athlete derives from doping. b has support $[\underline{b}, \overline{b}]$, pdf $f(b)$ and is independent of ϵ. It is assumed that $(1 - \theta n)(1 - \rho n)^{-1} \in [\underline{b}, \overline{b}]$ for all $n \in [0, 1]$. When a specific distribution must be specified, it is assumed that $b \sim N(\mu, \sigma^2)$. This can be justified by the central limit theorem since b can be viewed as the athlete's average benefit from doping over his career.

Athletes know their individual b but WADA only knows the underlying distribution from which b is drawn. If an athlete produces a performance of quality $a + \epsilon$, his lifetime income is $g(a + \epsilon)$. If he produces a performance of $ab + \epsilon$, his lifetime income is $g(ab + \epsilon)$. Reservation earnings are normalised to zero. For tractability it is assumed that $g(x) = zx$, where $z \geq 0$. The key elements of $g(x)$ are that it is increasing in x and is additively separable.[1] Since ϵ is mean zero, it does not enter into the analysis of the athlete's decision regarding whether or not to dope or into the expected social welfare function. The role of this stochastic element of performance is to ensure that consumers cannot perfectly infer from performance whether or not the athlete took drugs.

For a number of reasons, an athlete who takes drugs may not test positive, and an athlete who does not take drugs may test positive. ρ is the probability of a true positive, i.e. the probability that an athlete who takes drugs tests positive. An athlete who takes drugs may not test positive if no test exists for the drug the athlete is taking, or the athlete is able to mask the substance. θ is the probability of a false positive. False positives may arise due to laboratory errors or in rare cases where an athlete has an abnormally high naturally

1 In particular, if $g(x) = y + zx$ the results do not significantly change but the analysis becomes more complicated since y and z enter into the SWF.

occurring level of a substance. Given modern testing procedures it is reasonable to assume that $\rho > \theta$. Therefore $0 < \rho - \theta \leq 1$. If an athlete is banned from competing he receives only his reservation earnings for the period of the ban.

The timing of the model is as follows. First WADA announces the punishment for taking drugs. Athletes then privately learn their gain from taking drugs and subsequently make a one off decision as to whether or not to dope.[2] Finally, athletes are tested once prior to the start of their careers and the result of this test determines whether or not they are punished.[3] The anti-doping agency is assumed to be able to pre-commit to the optimal value of n, for example, because it is enshrined in law. If this were not the case, then given the set-up of the social welfare function, WADA would always want to deviate and set $n = 0$ after an athlete has decided whether or not to take drugs.

Since this model only involves one athlete, for this model to provide sensible results, it must be assumed that an athlete makes his decision about whether or not to take drugs independently of the choice of any other athlete. This approach has been taken previously in the literature and the decision by athletes in this model can be seen as a reduced form of the interaction between athletes. Athletes are assumed to be expected utility maximisers and consequently an athlete takes drugs if his expected utility from doing so exceeds that from not:

$$
\begin{aligned}
E\Big[(1-\rho)g\big(ab+\epsilon\big)+\rho(1-n)g\big(ab+\epsilon\big)\Big] &\geq E\big[(1-\theta)g(a+\epsilon)+\theta(1-n)g(a+\epsilon)\big] \\
\Rightarrow E\Big[(1-\rho n)g\big(ab+\epsilon\big)\Big] &\geq E\big[(1-\theta n)g(a+\epsilon)\big] \\
\Rightarrow b &\geq \frac{1-\theta n}{1-\rho n} \equiv \beta(n)
\end{aligned}
$$

2 This is clearly a simplifying assumption. As an athlete nears retirement, he may have a greater incentive to take drugs since he can only be banned for a shorter period of time. In this model, it is as if when an athlete nears retirement, his potential monetary gain from taking drugs is exactly offset by the increased cost to his reputation of being caught and the public believing that all of the athlete's results throughout his entire career were due to drug taking.

3 Alternatively, the athlete could be tested with a given probability and this probability could be incorporated into ρ and θ. The key assumption is that testing is independent of the punishment chosen by WADA. This would be true if the organisation responsible for carrying out testing simply maximised its number of tests subject to a budget constraint.

There is therefore a cutoff value of b such that the athlete takes drugs if and only if $b \geq \beta(n)$. This is increasing in n and consequently the right hand side of this inequality is smallest when $n = 0$. Therefore an athlete will only ever take drugs if his $b \geq 1$. Intuitively, even if there is no punishment for testing positive, an athlete will only wish to dope if his performance improves when he does so.

The participation constraint for an athlete who chooses not to take drugs is:

$$\begin{aligned} (1-\theta n)E\big[g(a+\epsilon)\big] &\geq 0 \\ \Rightarrow a &\geq 0 \end{aligned}$$

The participation constraint for an athlete who chooses to take drugs is:

$$\begin{aligned} (1-\rho n)E\Big[g\big(ab+\epsilon\big)\Big] &\geq 0 \\ \Rightarrow a &\geq 0 \end{aligned}$$

since $0 < \rho \leq 1$ and b is always at least 1 for an athlete who chooses to take drugs. Whether an athlete takes drugs or not, the participation constraint therefore reduces to $a \geq 0$. From now on it is assumed that this is the case and consequently the participation constraints do not impact on the subsequent analysis. This is because when the punishment is a ban, an athlete can never be fined more than his lifetime income.

It is assumed that society values higher quality performances. There is strong evidence for this since spectators and television channels are willing to pay far more to watch and broadcast the Olympic Games than a school level competition. If an athlete does not take drugs and does not test positive then society receives benefit a from the athlete's career. However, if the athlete tests positive, he cannot compete for proportion n of his career and society derives no benefit from the athlete for this period. Therefore the expected benefit which society obtains from an athlete who does not dope is $(1-\theta n)a$. Correspondingly if an athlete chooses to take drugs and has a gain from drug taking of b, society's expected benefit from this athlete is $(1-\rho n)ab$. However, if an athlete tests positive then society suffers a lump-sum loss L, regardless of whether or not the athlete actually took drugs. This reflects the loss society suffers when it emerges that a role model has taken drugs.

The expected social welfare function which WADA seeks to maximise is:

$$W(n) \quad = \quad \int_{\underline{b}}^{\beta(n)} \left[(1-\theta n)a - \theta L\right] f(b)\mathrm{d}b + \int_{\beta(n)}^{\overline{b}} \left[(1-\rho n)ab - \rho L\right] f(b)\mathrm{d}b$$

$$s.t.\ 0 \leq n \leq 1$$

By assumption, $\beta(n)$ lies within the support of b for all possible parameter values.

The first integral of $W(n)$ refers to the social welfare derived from an athlete who does not take drugs, weighted by the probability that the athlete under consideration will choose not to take drugs. The second integral shows the expected social welfare derived from an athlete who dopes, weighted by the probability that the athlete will choose to dope. The first term of each integrand refers to the benefit society derives from the performance of the athlete, adjusted to take into account the probability that the athlete tests positive and is therefore banned for proportion n of his career. The second term refers to the expected social cost as a result of society suffering a loss if the athlete tests positive.

The second integral shows that society benefits from the performance of an athlete who takes drugs, except if the athlete is caught, in which case society receives no benefit for the period of the ban. The intuition for this is that society cannot identify from performance alone whether or not the athlete cheated. In this chapter, consumers do not enter into the game as players and it is simply assumed that society values higher performances. This assumption is relaxed in chapter 3 where consumers are fully rational players and engage in Bayesian updating.

An additional important assumption is that if the athlete tests positive, society values the athlete's performance for the proportion of his career for which he competes. This reflects the idea that in a dynamic setting, consumers enjoy watching high performances until it is revealed that the athlete failed a drug test, at which point the athlete is banned and society no longer derives welfare from the athlete. In the set-up of this game, the assumption can be justified if the ban is timed such that the athlete retires at the end of his ban and the positive test is only revealed at the start of the ban. Alternatively, it holds if society does not care that the athlete doped in terms of valuing his performance. If instead test results were revealed at the start of the game and society does not value the performance of an athlete who tested positive even if he competes, n would not enter into the integrands and the only effect of a change in n would be to decrease the proportion of athletes taking drugs. An increase in n then unambiguously increases social welfare and the optimal punishment is always to ban the athlete for life if he tests positive.

Given this assumption there are two effects of a change in n. Firstly, a change in n changes the proportion of athletes who choose to dope. When n increases, $\beta(n)$ also increases. Since an athlete only cheats if he has $b \geq \beta$, an increase in β means that fewer athlete types will choose to cheat. Therefore by changing the limits of integration, a change in n effectively alters the relative weight attached to the two terms in the social welfare function. Secondly, an increase in n increases the punishment for an athlete who tests positive. This harms social welfare since it decreases the proportion of an athlete's career for which society can expect to benefit from that athlete's performance.

2.3.2 *First and second order conditions*

Let $(1-\theta n)/(1-\rho n) \equiv \beta(n)$. Then:

$$\frac{d\beta(n)}{dn} = \frac{\rho-\theta}{(1-\rho n)^2}$$

Consequently, $\frac{d\beta(n)}{dn} > 0$ and when the punishment increases, fewer types of athlete choose to take drugs.

The first order condition for WADA's choice of n is:

$$\underbrace{-\theta a F(\beta^*) - \rho a\big(1-F(\beta^*)\big)\mathrm{E}\Big[b \mid b \geq \beta(n^*)\Big]}_{\textit{Punishment Channel}} + \underbrace{\left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 L f(\beta^*)}_{\textit{Proportion Channel}} = 0 \quad (2.1)$$

where $\beta^* = (1 - \theta n^*)/(1 - \rho n^*)$. The derivation of this is given in appendix A.1.

When contemplating the optimal n to implement, WADA must trade off two counteracting forces. Firstly, when n increases while the proportion of athletes taking drugs is held constant, there is a social loss since any athlete who tests positive will be banned for a longer proportion of his career. This is captured by the first term of equation 2.1. An athlete who does not take drugs tests positive with probability θ and is banned for proportion n^* of his career. Therefore when n increases, the expected benefit society receives from an athlete who does not take drugs declines by θa. The probability that the athlete under consideration is a type who does not take drugs is $F(\beta^*)$ since an athlete does not take drugs if his $b < \beta$. If an athlete takes drugs, the expected lifetime benefit he generates decreases by $\rho a \mathrm{E}\big[b \mid b \geq \beta(n^*)\big]$ when n increases. $1 - F(\beta^*)$ is the probability that the athlete dopes. Therefore the first term of equation 2.1 is the expected change in social welfare due to an athlete being banned from competing for longer when n increases, taking into account the probability with which the athlete under consideration dopes. This is termed the *punishment channel,* and an increase in n always affects social welfare negatively via this channel.

However, there is a second effect since the proportion of athlete types taking drugs also decreases when n increases. This is termed the *proportion channel* and it is captured by the second term of equation 2.1. When fewer types of athlete choose to take drugs, the expected social cost of drug taking decreases since the probability of an athlete who takes drugs testing positive (ρ) exceeds the probability for an athlete who does not take drugs (θ). On the other hand, when fewer types of athlete dope, this decreases the expected social benefit from the athlete's performance if the athlete does not test positive. However, once the decreased risk of the athlete testing positive is taken into account, a marginal increase in n has no impact on the expected benefit which society receives from performance over the athlete's lifetime. Therefore the *proportion channel* always has a positive impact on social welfare.

At the optimum, the decline in social welfare via the *punishment channel* is exactly offset by the increase in social welfare via the *proportion channel*. If L is sufficiently large, the positive effect of an increase in n via the *proportion channel* will always outweigh the negative effect via the *punishment channel* and $n^* = 1$. On the other hand, if baseline performance, a, or the expected improvement in performance when an athlete takes drugs, $\mathrm{E}\big[b \mid b \geq \beta(n^*)\big]$, are sufficiently large, the negative effect which the *punishment channel* has on social welfare will outweigh the positive effect from the *proportion channel.* In this case it is optimal not to ban an athlete at all even if he tests positive.

The second order condition is:

$$\frac{(\rho-\theta)^2}{(1-\rho n^*)^3} f(\beta^*)\left[2\rho L + \frac{\rho-\theta}{1-\rho n^*}\frac{f'(\beta^*)}{f(\beta)} L + a\right] \quad \leq \quad 0$$

The derivation of this is given in appendix A.2. A necessary condition for the second order condition to be satisfied is $f'(\beta^*) < 0$.

2.3.2.1 *b normally distributed*

If b is normally distributed, the SWF can be rewritten as:

$$W(n) = (1-\theta n)a\Phi(\alpha) + (\rho-\theta)L\Phi(\alpha) - \rho L + (1-\rho n)a\sigma \int_{\alpha}^{\infty}\left(z + \frac{\mu}{\sigma}\right)\phi(z)\mathrm{d}z$$

where $\alpha = (\beta - \mu)/\sigma$ and $z = (b-\mu)/\sigma$.

Using this formulation, the first order condition for WADA's choice of n is:

$$\underbrace{-\theta a\Phi(\alpha^*) - \rho a\big(1-\Phi(\alpha^*)\big)\mathrm{E}\Big[b \mid b \geq \beta(n^*)\Big]}_{\textit{Punishment Channel}} \;+ \underbrace{\left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 L\sigma^{-1}\phi(\alpha^*)}_{\textit{Proportion Channel}} = \;\; 0 \quad (2.2)$$

where $\alpha^* = (\beta^* - \mu)/\sigma$ and $\beta^* = (1-\theta n^*)/(1-\rho n^*)$.

The second order condition is:

$$\frac{(\rho-\theta)^2}{\sigma(1-\rho n^*)^3}\phi(\alpha^*)\left[2\rho L + a - \left(\frac{\rho-\theta}{1-\rho n^*}\right)\left(\frac{\beta^*-\mu}{\sigma^2}\right)L\right] \quad \leq \quad 0$$

A necessary condition for the second order condition to be satisfied is $\beta^* > \mu$.

2.3.3 *Comparative statics*

The comparative statics derived below hold as long as an interior solution exists. The derivatives are derived in appendix A.3. Since the second derivative of social welfare is negative provided we have an interior solution, if $\partial W'(n)/\partial x$ is positive, when x increases this implies n^* should also increase.

The derivative of $W'(n)$ with respect to L is:

$$\frac{\partial W'(n)}{\partial L} = \left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 f(\beta^*) > 0$$

This is unambiguously positive which is intuitive. An increase in L increases the positive effect which the *proportion channel* has on social welfare while leaving the *punishment channel* unchanged. When the social cost of drug taking increases, the optimal punishment also increases in order to deter more types of athlete from taking drugs.

The derivative of $W'(n)$ with respect to a is:

$$\frac{\partial W'(n)}{\partial a} = -\theta F(\beta^*) - \rho\big(1 - F(\beta^*)\big)\mathrm{E}\big[b \mid b \geq \beta(n^*)\big] < 0$$

This is unambiguously negative meaning that when a increases the optimal punishment decreases. This effect operates entirely via the *punishment channel.* When a increases, any given ban decreases the benefit which society obtains from an athlete's career by a greater amount. This implies the punishment should decrease in order to offset this effect.

Further analysis can only be carried out if a distribution for b is specified. Therefore for the remainder of the chapter it is assumed that b is normally distributed. As was discussed above, this can be justified by the central limit theorem. It is then possible to analyse the impact of changes in μ and σ on the optimal punishment. The derivative of the first order condition with respect to μ is:

$$\frac{\partial W'(n)}{\partial \mu} = \underbrace{-\frac{\rho-\theta}{\sigma(1-\rho n^*)}a\phi(\alpha^*) - \rho a\big[1-\Phi(\alpha^*)\big]}_{\text{Change in effect of } \textit{punishment channel}} \underbrace{+\left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 \frac{\alpha^*}{\sigma^2} L\phi(\alpha^*)}_{\text{Change in effect of } \textit{proportion channel}}$$

where $\alpha^* = (\beta^* - \mu)/\sigma$.

From the second order condition we know $\beta^* > \mu \Rightarrow \alpha^* > 0$. Therefore the sign of this is ambiguous because an increase in μ strengthens both the negative effect of the *punishment channel* and the positive effect of the *proportion channel*. Which of these effects dominates depends on the exact parameter values. When μ increases, an increase in n decreases social welfare to a greater extent via the *punishment channel* because average performance with drug taking increases. This means that for a given ban, if an athlete who is doping tests positive, the loss of social benefit is now greater. On the other hand, when $\beta^* > \mu$, if μ increases, for a given increase in n, the proportion of athletes doping decreases by a greater extent. Therefore the overall relationship between μ and n^* is ambiguous. If L is sufficiently large however, the effect via the *proportion channel* will dominate and $\mathrm{d}n^*/\mathrm{d}\mu > 0$.

The derivative of the first order condition with respect to σ is:

$$\frac{\partial W'(n)}{\partial \sigma} = \underbrace{-\frac{\rho - \theta}{\sigma(1 - \rho n^*)} a\alpha^* \phi(\alpha^*) - \rho a \phi(\alpha^*)}_{\text{Change in effect of } \textit{punishment channel}} \underbrace{+ \left(\frac{\rho - \theta}{(1 - \rho n^*)\sigma}\right)^2 (\alpha^2 - 1) L\phi(\alpha^*)}_{\text{Change in effect of } \textit{proportion channel}}$$

The sign of this is ambiguous. A change in σ strengthens the negative effect of the *punishment channel* and may either strengthen or weaken the positive effect of the *proportion channel*. When σ increases and $\beta^* > \mu$, average performance with drug taking increases. Therefore, a given ban now has a more negative effect on social welfare via the *punishment channel* since the social benefit from the career of an athlete who dopes increases. However, the effect of a change in σ on the *proportion channel* is ambiguous. If $\alpha^2 > 1$, then as σ increases, the proportion of athletes taking drugs decreases by a greater extent for a given increase in n. This increases the beneficial effect which an increase in n has on social welfare via the *proportion channel*. Therefore if $\alpha^2 > 1$, the effect of a change in σ

on n^* is ambiguous. However, if $\alpha^2 < 1$ then an increase in σ unambiguously decreases the optimal length of ban.

Since for two key parameters of interest, μ and σ, the comparative statics are ambiguous, simulations are carried out to better understand the effect of changes in these parameters.

2.4 CALIBRATION AND SIMULATIONS

2.4.1 *Calibration*

2.4.1.1 *Distribution of b*

For the simulations it is assumed that b is normally distributed with mean μ and standard deviation σ. The results are qualitatively similar if it is instead assumed that b is log normally distributed. With the log normal distribution, as σ^2 increases not only the variance, but also the skewness of the distribution increases.

When a value of μ must be specified it will be assumed that $\mu = 1$. This is intuitive since it means the median of b is also 1. Therefore when the punishment is equal to 0, so that $\beta(0) = 1$, half of the population of athletes will not take drugs, while half will. Consequently, the assumption that $\mu = 1$ can be translated into an assumption that half of the population have sufficiently high morals that even if they would not be punished for taking drugs, they would not wish to do so.

There is no intuitive value to assign to σ. Since b denotes the gain an athlete obtains from taking drugs, σ^2 determines the variance in the gain an athlete derives from taking drugs. Little is known about the variance in performance improvement in response to doping. However, b can also be thought of as determining the reward an athlete derives from an improvement in performance. Since more is known about the distribution of earnings in sport than the distribution of gains from drug taking, it is simpler to think about σ^2 in this context. The distribution of earnings is likely to vary widely across sports. Scully (1995) discusses potential factors which are likely to result in a more or less equal distribution of earnings among athletes and then tests these predictions empirically. Scully predicts that earnings inequality in individual sports is likely to be greater than in team sports because in team sports, opponents are often able to exploit the weakest link. When team sports are compared to golf, there is found to be less earnings inequality in team sports than in individual sports. Scully also suggests that sports with more competitions are likely to have greater income inequality as a result of reduced randomness in outcomes. It is therefore likely that the appropriate value of σ will differ across sports with team sports and sports with fewer competitions having a lower value of σ.

2.4.1.2 *a*

a is the baseline benefit society derives from the athlete's career if he does not take drugs. Consequently a will vary across sports depending on the average number of years in an athlete's career and the extent to which society values performances in the sport under consideration. It is not necessary to estimate a in order to carry out the simulations. This is because the value of n which maximises $W(n)$ will also maximise:

$$\frac{W(n)}{a} = \int_{-\infty}^{\beta(n)} \left[(1-\theta n) - \theta\tau\right] f(b)\mathrm{d}b + \int_{\beta(n)}^{\infty} \left[(1-\rho n)b - \rho\tau\right] f(b)\mathrm{d}b$$

where $\tau = L/a$ is the ratio of the social cost of doping to the benefit derived from the baseline performance.

2.4.1.3 *Probability of a false positive, θ*

In the model proposed here, the athlete makes a one-off decision about whether or not to take drugs and then faces a different probability of being caught and banned depending on whether or not he has taken drugs. In order for the simulations to be realistic, the probabilities of false and true positives must reflect the athlete's lifetime probability of testing positive i.e. θ should be an athlete's probability of testing positive once in his career if the athlete has not taken drugs. This is potentially different from the probability of an athlete testing falsely positive in any one given drug test.

False positives do not generally occur due to technical laboratory errors, but rather because in rare cases athletes may have a naturally occurring level of a banned substance which is above that deemed to mark the threshold at which an athlete is declared to have taken drugs. This means that the probability of a false positive occurring in an athlete's lifetime and the probability of it occurring in a given test should be very similar since an athlete who has a natural level of the substance which is above the threshold will frequently test positive. WADA has stated that for the test for human growth hormone, the false positive rate is at least 1 in 10,000 (Macur, 2011). This is also the figure for a false positive which WADA stated with regard to the test for EPO in 2000 (WADA, 2000). Therefore in the simulations carried out below, it is assumed that $\theta = 0.0001$.

2.4.1.4 *Probability of a true positive, ρ*

It is unlikely to be possible to compute a value for the probability that an athlete who takes drugs tests positive during his career. The key probability here is not so much the probability that a lab incorrectly labels a sample negative (although this is clearly an issue), but rather the probability of an athlete being tested when he has taken drugs, the probability of the anti-doping agency having already developed a test for the substance the athlete is taking and the probability that the athlete does not manage to evade the test or mask the evidence that he has taken drugs. Even the probability that an athlete is tested out of competition cannot be calculated since anti-doping agencies do not publish this statistic. Clearly the probability that an athlete who has taken drugs tests positive once in his career is not 1, as evidenced by the recent announcement by Lance Armstrong that he took drugs in every Tour de France that he won (BBC, 2013). Given the uncertainty surrounding this figure, some simulations will need to be conducted using a range of values for ρ. When a value of ρ must be specified in the simulations, 0.9 will be used. This captures the intuition that testing is by no means perfect, but nevertheless, the majority of athletes who take drugs are caught at some point in their lifetime.

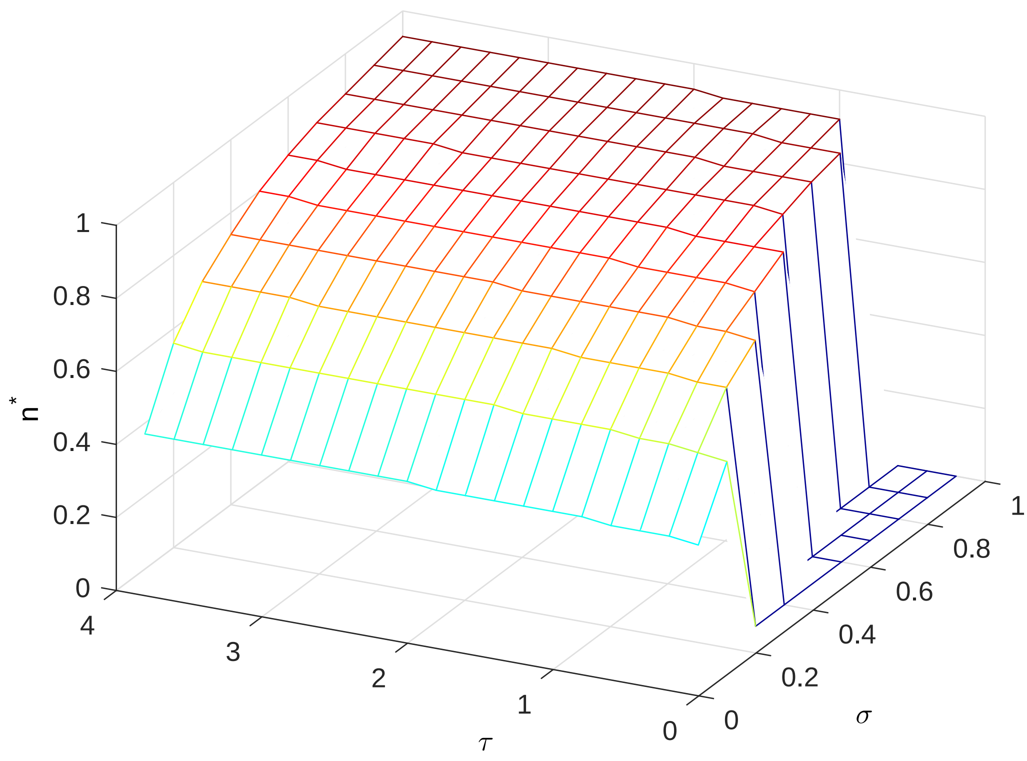

Figure 2.1: Optimal value of n as τ and σ vary

2.4.2 *Simulations*

2.4.2.1 *Optimal value of n as τ and σ vary*

This simulation was carried out using $\rho = 0.9$. The results of this simulation are shown in figure 2.1.

An increase in τ $(= L/a)$ corresponds to either an increase in L, the social cost of doping, or a decrease in a, baseline performance without doping. The comparative statics showed that when L increases or a decreases, the optimal punishment should increase. This is confirmed by the results of the simulation shown in figure 2.1. However, when τ is low we have a corner solution in which $n^* = 0$ and an athlete who tests positive is not banned from competing for any length of time. This is because when τ is low the *proportion channel* has little positive effect on social welfare since there is little benefit from deterring athletes from taking drugs. The negative effect of the *punishment channel* then dominates and this implies that n^* should be set as low as possible.

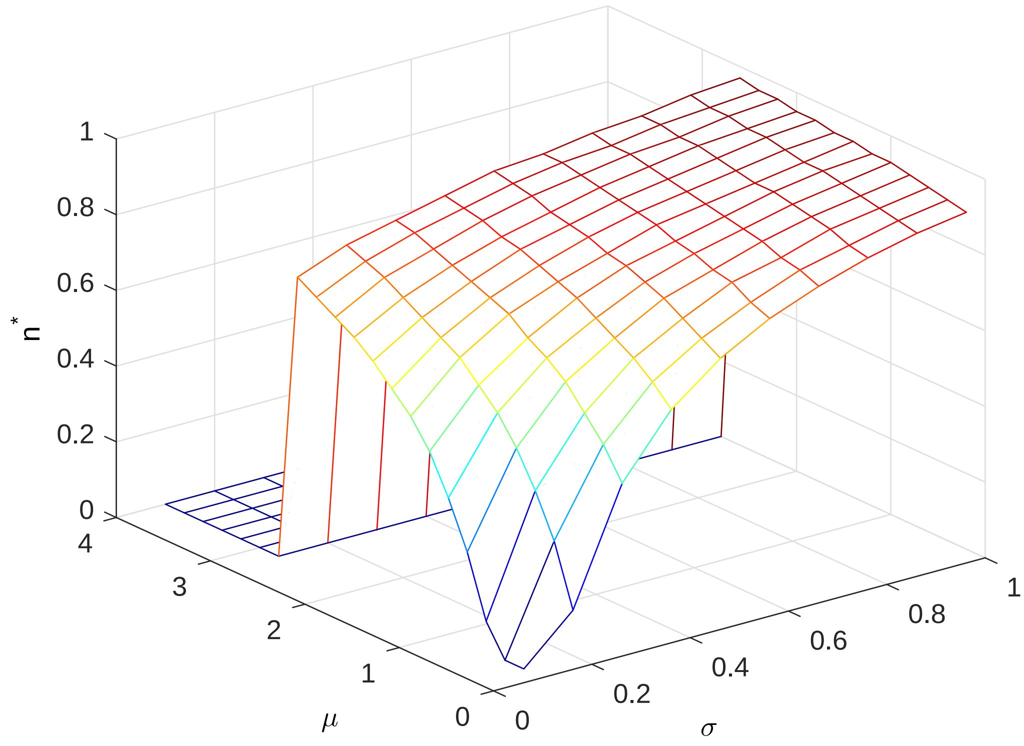

Figure 2.2: Optimal value of n as μ and σ vary

In all of the simulations carried out, provided we do not have a corner solution, as σ increases, the optimal punishment also increases. In the comparative statics the effect of an increase in σ was ambiguous. This was because while an increase in σ strengthened the negative effect due to the *punishment channel* it could also strengthen the positive effect of an increase in n^* via the *proportion channel*. The simulations therefore suggest that the positive effect from the *proportion channel* outweighs the negative effect from the *punishment channel*.

2.4.2.2 *Optimal value of n as μ and σ vary*

This simulation was carried out using $\rho = 0.9$ and $\tau = 2$. The results of this simulation are shown in figure 2.2.

Once again, as σ increases n^* increases provided we do not have a corner solution. At first as μ increases, n^* increases. Over the range of μ for which we have an interior solution, this implies that the effect of a change in μ on the *proportion channel* domin-

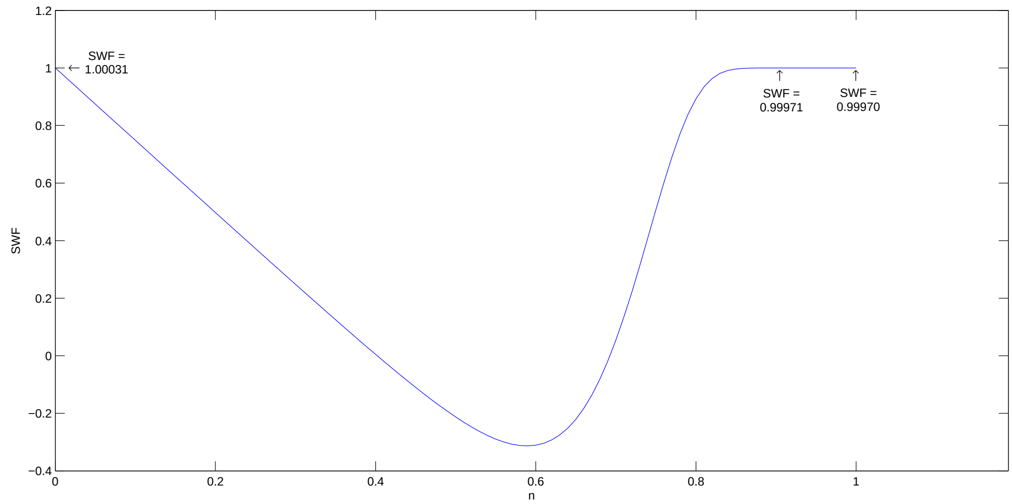

Figure 2.3: SWF as n varies

ates the effect on the *punishment channel*. When μ increases, the benefit of increasing n and having fewer types of athletes choose to dope exceeds the benefit of decreasing n and benefiting from the performances of athletes for a greater proportion of their careers. However, once a certain point is reached there is a jump in the optimal punishment to zero. The reason for this can be seen by examining a plot of the SWF versus n when $\mu = 2.8$ and $\sigma = 0.5$. This is shown in figure 2.3.

The SWF reaches a local maximum at $n = 0.9$. The key change which occurs when the optimal n^* jumps down to zero is not that the shape of the SWF fundamentally changes but rather that the value of the SWF when $n = 0$ increases so that it exceeds the value at the interior maximum. Letting β_I denote the value of β when $n = n_I$ is the local interior maximum and noting that when $n = 0$, $\beta = 1$, we have a corner solution with $n^* = 0$ if SWF(0)$\geq$SWF(n_I). This is true if:

$$\begin{aligned}\mu \;\; \geq \;\; & \frac{1}{1-\Phi\left(\frac{1-\mu}{\sigma}\right)-(1-\rho n_I)\left[1-\Phi\left(\frac{\beta_I-\mu}{\sigma}\right)\right]} \times \\ & \left[(1-\theta n_I)\Phi\left(\frac{\beta_I-\mu}{\sigma}\right)-\Phi\left(\frac{1-\mu}{\sigma}\right)+\sigma\left[(1-\rho n_I)\phi\left(\frac{\beta_I-\mu}{\sigma}\right)-\phi\left(\frac{1-\mu}{\sigma}\right)\right]\right. \\ & \left.+(\rho-\theta)\tau\left[\Phi\left(\frac{\beta_I-\mu}{\sigma}\right)-\Phi\left(\frac{1-\mu}{\sigma}\right)\right]\right] \equiv \underline{\mu}\end{aligned}$$

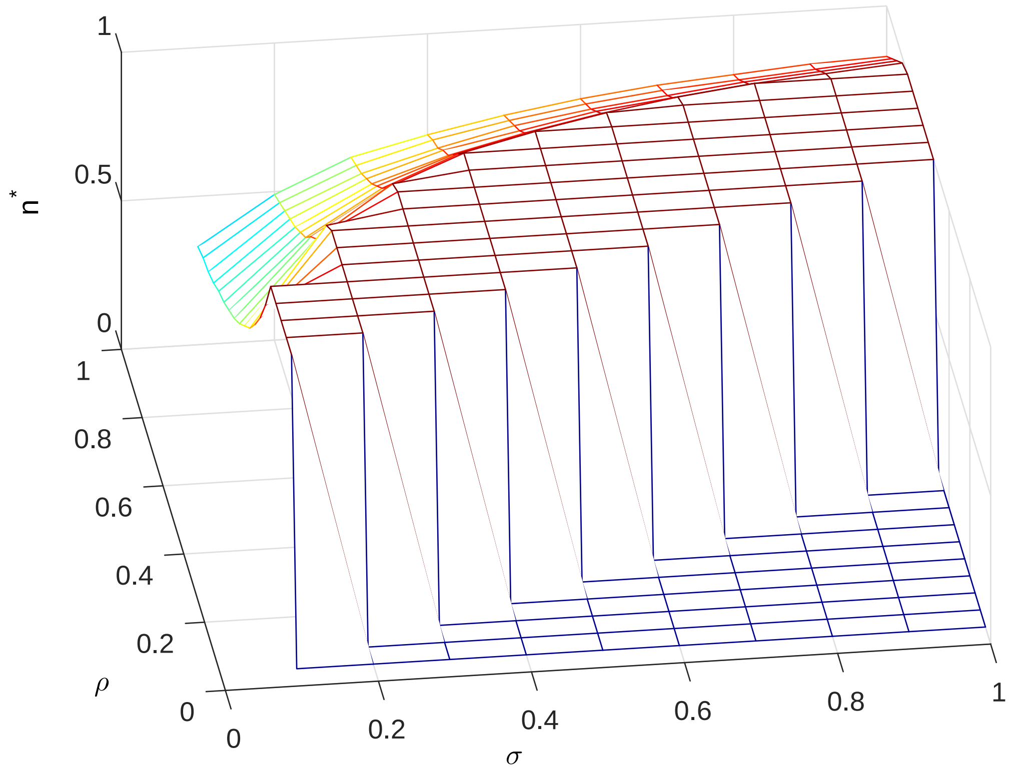

Figure 2.4: Optimal value of n as ρ and σ vary

Therefore when $\mu \geq \underline{\mu}$ we have a corner solution with $n^* = 0$.

2.4.2.3 *Optimal value of n as ρ and σ vary*

This simulation was carried out using $\tau = 2$. The results of this simulation are shown in figure 2.4.

For large ranges of ρ and σ there is a corner solution. This is because when ρ is relatively small, the SWF is convex over all possible values of n and we have a corner solution with either $n^* = 0$ or $n^* = 1$. For low values of ρ, the optimal punishment is zero since there is little risk of an athlete testing positive and therefore little risk of society incurring the social cost of doping even if the athlete takes drugs. However, for higher values of ρ but values where the SWF is still globally convex, the positive effect of an increase in n via the *proportion channel* always outweighs the negative effect via the *punishment channel* and instead the optimal punishment is a lifetime ban. Once ρ is sufficiently large that the second order condition is satisfied, n^* decreases as ρ increases.

2.5 DISCUSSION

The simulations above suggest that the optimal punishment is highly dependent on the values of τ, σ, μ and ρ. As is discussed below, these parameters are all likely to differ across sports. This analysis consequently supports the argument that social welfare could potentially be increased if WADA were to relax its requirement that sanctions be standardised across sports. The discussion below highlights some of the considerations which should be taken into account if WADA were to change its policy.

Firstly, the results unambiguously suggest that as τ increases the optimal punishment should also increase. $\tau = \frac{L}{a}$ where L is the social cost of doping and a is the benefit which society derives from the career of an athlete who does not take drugs. *Ceteris paribus*, τ will be larger the larger is L or the smaller is a. A higher value of L results in a greater value of τ implying that the optimal punishment will be higher. In assessing the relative size of L across different sports, a fundamental question is what precisely the social cost of doping captures. Most would not dispute that there is a social cost when an athlete tests positive, but it is harder to determine the precise causes of this cost. Clearly at the most basic level, society is faced with potential health costs for an athlete who takes drugs. However, much of the social cost of doping is intangible and when discussing the ills of doping, commentators frequently focus on aspects such as the moral cost, since athletes are often role models for both children and society more generally. If such intangible costs are deemed to be important, their magnitude is likely to be greater in sports which attract a higher profile, especially among children. This suggests that sports such as football are likely to have a higher social cost of doping than sports such as curling.

As a increases, τ decreases and the optimal punishment also decreases. a can be seen as yearly benefit from an athlete competing without doping multiplied by the average number of years in an athlete's career. If sports with shorter career lengths have the same yearly benefit from performance without drug taking as sports with a longer career length, *ceteris paribus* it will be optimal for sports with shorter career lengths to have a higher n^* and for those with longer career lengths to have a lower n^*. This may in fact mean that sports with different career lengths have a similar length of ban. If career length is shorter but the yearly benefit from performance is the same, society derives less benefit from this sport over the athlete's lifetime and therefore there is less loss to banning an athlete for a greater proportion of his career.

Another dimension over which sports differ is with respect to their value of μ. μ is the mean multiplicative gain which an athlete derives from taking drugs. μ could therefore vary for two main reasons. The first is that drugs are likely to be more useful, and subsequently provide a greater performance benefit, in some sports than others. For example, drugs tend to be more useful in sports which are predominantly strength, speed or endurance based as opposed to sports in which skill and tactics are important. Therefore we might expect sports such as cycling, where EPO and blood doping are known to be highly effective, to have a higher value of μ than sports such as sailing. The second reason why μ could vary is that cultures of doping are believed to be more pervasive in some sports than others. In conjunction with σ, μ determines $\Phi(\frac{1-\mu}{\sigma})$, the proportion of athletes who choose not to dope even if $n^* = 0$. An increase in μ causes $\Phi(\frac{1-\mu}{\sigma})$ to decrease meaning a higher proportion of athletes will choose to dope when $n = 0$. Consequently sports in which doping is viewed as more acceptable by athletes can be seen as having a higher μ. For example, the Cycling Independent Reform Commission (2015) has found that a 'culture of doping' has traditionally been present in cycling. Meanwhile, Strulik (2012) analyses how peer-approval of doping can influence an athlete's decision regarding whether or not to dope. He shows that doping is more likely to emerge in sports where there is less stigma surrounding doping among peers and where group cohesion is stronger.

The simulations show that as μ increases, at first the optimal punishment also increases. However, once a threshold value of μ is reached, the optimal punishment instead drops to zero. This suggests that in sports where the gain from taking drugs is only moderate, or where a substantial proportion of athletes are willing not to dope even if they would not be punished for doing so, it is optimal to ban athletes for a non-negligible proportion of their career. This reduces the expected social costs of doping and the lost benefit from performance when an athlete is banned is not overly large. On the other hand, if there are substantial performance benefits from doping or a doping culture is endemic in a sport, the simulations suggest it is better not to attempt to deter athletes from doping. Attempts at deterrence then have little impact on the proportion of athletes choosing to dope but are costly to society in terms of expected performance foregone. This supports the viewpoint that if doping is already endemic in a sport we should potentially accept it rather than trying to fight it.

Sports also differ with respect to their value of σ. Over ranges of parameter values for which we have an interior solution, as σ increases the optimal length of ban also increases. Most directly, σ^2 is the variance in the gain which athletes obtain from taking drugs. However, as was discussed above, σ^2 can also be thought of as determining the distribution of rewards athletes receive as performance varies. As was discussed in section 2.4.1.1, Scully (1995) has found that team sports and sports with fewer competitions are likely to have lower variance of earnings, corresponding to a lower value of σ, than individual sports and those with many competitions. This implies that all else equal, sports such as hockey should have a lower punishment than sports such as athletics and swimming.

Finally, if ρ differs across sports, this suggests a further reason why optimal sanctions should not be uniform across sports. ρ is the probability that an athlete who takes drugs tests positive. An athlete may not test positive despite taking drugs if he is not tested, or if he takes a drug for which a test is yet to be developed by the anti-doping agency. Consequently, sports with less funding for testing, or those where new drugs are being developed, may have a lower value of ρ. The simulations showed that when there is an interior solution, as ρ increases, n^* declines. This suggests that in sports where there is deemed to be a high risk of athletes not testing positive if they take drugs, the punishment for testing positive should be harsher.

2.6 CONCLUSION

Previous papers have often assumed that it is optimal to implement a no-doping equilibrium. However, by accounting for the fact that the sanction in sport is a ban, not a lump sum fine, this chapter shows that maximal deterrence is not always optimal. The chapter then examines how the optimal punishment should vary with the characteristics of different sports. It has been shown that the optimal punishment varies with the popularity of the sport, the prevalence of doping, the variance of earnings and the testing regime.

There are several possible extensions to the model developed here. Firstly, it would be desirable to develop a dynamic model whereby athletes can change their decision about whether or not to take drugs and can also be tested throughout their careers so that the ban which WADA can implement becomes shorter as athletes near retirement. Secondly, it would be interesting to endogenise ρ and θ for the anti-doping agency so that the testing regime can be improved, but this is costly. Finally, the game in this chapter did not model interactions between athletes in the decision by the athlete of whether or not to take drugs. Although the decision by athletes in this model can be seen as a reduced form of such an interaction, it could nevertheless be insightful to model this interaction explicitly.

3 INCENTIVISING ANTI-DOPING AGENCIES TO TEST ATHLETES

3.1 INTRODUCTION

In 1998 Lance Armstrong returned to competitive cycling as a cancer-survivor. For the International Cycling Union (UCI), his return could not have come at a better time. Armstrong was a hero who appeared to herald a new era of drugs free cycling. His seemingly miraculous come-back brought disillusioned fans and sponsors alike back into the sport. However, during the 1999 Tour de France, a crack appeared in the facade: Armstrong tested positive for the banned steroid Cortisone (USADA, 2012). The reputation of cycling would have been irrevocably damaged if Armstrong had been revealed to be a fraud. According to Armstrong, the UCI therefore chose to accept a backdated prescription for a saddle sore cream containing the banned substance, thereby covering up evidence of his positive test (Huffington Post, 2013). Meanwhile, in a particularly self-critical report tellingly entitled *'Lack of Effectiveness of Testing Programs' (WADA, 2013*b*),* the World Anti-Doping Agency (WADA) states there are 'inherent conflicts of interest' within the organisations responsible for testing athletes. This has been seen most recently with the revelations that the Russian Anti-Doping Agency systematically helped Russian athletes to avoid testing positive (Pound et al., 2015).

This chapter examines situations in which those responsible for testing athletes may lack incentives to uncover cheating. This could occur for a variety of reasons. For example, the governing body of a sport may benefit when athletes perform at a higher level due to increased sponsorship and television revenues. On the other hand, the reputation of the body may suffer if an athlete tests positive and a positive test may result in consumers leaving the sport, where consumers could be either spectators or sponsors.

In the model presented in this chapter there are three players: an anti-doping agency (ADA), an athlete and consumers. There are two types of athlete: L and H. Type L produces a low performance if he does not cheat and a mid performance if he cheats. Type H produces a mid performance if he does not cheat and a high performance if he cheats. Therefore, if consumers only observe performance, they cannot distinguish between a type L athlete who cheated and a type H athlete who did not.

The athlete decides whether or not to cheat while the ADA decides the probability with which the athlete is tested. In reality, the ADA could take steps to ensure an athlete does not test positive in other ways besides failing to test the athlete. For example, the ADA could warn athletes about tests in advance or suppress positive tests once they have occurred. Therefore in this model a reduction in the testing probability can more broadly be interpreted as a reduction in the deterrence effect of testing as a result of the ADA taking such actions. Consumers compute the posterior probability that the athlete cheated based on the information available to them and depending on this probability decide whether or not to participate.

The current policy of anti-doping agencies is not to provide information to athletes or consumers about the frequency with which athletes are tested. This corresponds to the baseline case with simultaneous timing in the analysis below. The purpose of this policy is to ensure that athletes are not able to predict when they will be tested. However, while it is clearly advantageous for athletes not to be able to predict the occurrence of tests, the analysis shows that a no-cheating equilibrium only exists when the ADA pre-commits to the testing probability. Such a policy would be compatible with the occurrence of individual tests being unpredictable. When the ADA does not pre-commit to its testing probability, it always has an incentive to deviate and not test since the athlete cannot respond to such a deviation.

Nevertheless, making the testing frequency observable to the athlete is not sufficient to achieve a no-cheating equilibrium. In the baseline case, consumers are only able to observe the athlete's performance and positive tests. However, they do not observe the probability with which the athlete was tested and they cannot tell whether the absence of a positive test is due to the athlete not having been tested or having tested negative. In this case, even when timing is sequential, a no-cheating equilibrium does not exist. The difficulty is that when consumers cannot observe the testing probability, their actions cannot be conditional on the testing probability actually set by the ADA. The ADA then always has an incentive to deviate by not testing in which case a type L athlete will wish to cheat.

This chapter therefore examines two possible solutions to the problem of doping. In the first, consumers are able to observe the frequency with which the athlete was tested. In addition, consumers may or may not also be able to observe negative tests. The intuition for examining both of these possibilities is that an anti-doping agency may be willing to make information about testing probabilities public at the population level, but for privacy reasons would be unwilling to disclose information about negative tests on individual athletes. On the other hand though, if consumers do not observe negative tests, it may be hard for them to verify that the ADA is testing with the frequency it claims. When consumers observe the frequency of testing, a no-cheating equilibrium exists and is unique regardless of whether or not they observe negative tests, provided the fine the ADA must pay if an athlete tests positive is sufficiently high and consumers will only tolerate a low probability of cheating. However, it is shown that if the payoffs of the athlete are proportional to those of the ADA, a no-cheating equilibrium may exist with observability of negative tests when it does not exist when negative tests are not observable.

However, it may be unreasonable to assume that consumers can access and understand information on testing. Therefore, the second solution which is considered is to allow an external testing authority to randomly carry out tests and fine the ADA if it discovers athletes to be cheating who the ADA failed to identify. The ADA then cannot guarantee that an athlete who cheats will not test positive, unless it sets the testing probability sufficiently high that neither type of athlete wishes to cheat. If the fine the ADA must pay if the athlete fails the external's test is sufficiently high, the ADA will be incentivised to test with a high enough probability that neither type of athlete wishes to cheat.

The results of this chapter therefore suggest that a no-cheating equilibrium cannot be achieved given the current testing framework. Instead, the ADA must be incentivised to test with sufficient intensity that athletes are deterred from cheating. If consumers are to hold the testing agency to account, more transparency is required about the probability with which athletes are being tested. If anti-doping authorities are not willing to publish this information, it is instead necessary to allow an impartial testing authority to carry out independent random tests.

3.2 RELATED LITERATURE

The model presented here combines elements of a variety of different models in the literature. It draws on the literature on regulatory capture since those responsible for monitoring compliance act in the interest of the agent rather than society. Such models (for example Laffont and Martimort, 1999) commonly use a three tier set-up with a principal, supervisor and agent. The supervisor receives a signal about a parameter which affects the agent's profitability and he may choose whether or not to disclose this signal to the principal. In certain situations both the supervisor and the agent can improve their utility by engaging in a side-contract whereby the supervisor is bribed by the agent not to reveal the true signal to the principal. Models such as that developed in Kofman and Lawarree (1993) also allow for the possibility of an external auditor who can provide a check on the results of the internal auditor. Such models are the motivation for the examination of an external testing authority in my work.

While my model incorporates a conflict of interests for the supervisor (in my model the ADA), unlike in models of regulatory capture, in my model the supervisor is inherently aligned with the agent because the ADA benefits directly from the athlete's performance. Therefore the supervisor does not need to be bribed by the agent in order to set incentives such that the agent cheats. Furthermore, in my model the supervisor is responsible for designing the contract for the agent and by choosing the probability with which to test the agent, changes the agent's incentives as to whether or not to cheat.

In this respect, my model is closer to inspection games. In inspection games, an inspector decides whether or not to test an individual who must simultaneously decide whether or not to cheat. Such games have been applied extensively to situations such as arms control and compliance with environmental inspections (for example Avenhaus et al., 2002). However, unlike in inspection games, in my model the inspector is not benevolent. In a sporting context, Kirstein (2012) relaxes the assumption of perfectly informative tests to study a Bayesian enforcement game, but still assumes a benevolent inspector and has neither consumers nor an external testing authority. My model also allows for the possibility that false negatives arise in testing.

There is also a body of literature which focuses on how to incentivise truthful reporting when it is possible for an agent or monitor to manipulate reports. For example, Clausen (2013) has examined how the possibility of an agent being able to suppress some bad signals affects the design of optimal contracts in moral hazard problems. He finds that fraud can be deterred if more suspicious signals are rewarded less but the contract is more lenient on bad signals that could have been suppressed. Meanwhile, Rahman (2012) analyses how a monitor can be incentivised to detect deviations by those he is responsible for monitoring. He finds that compliance can be achieved without auditing by occasionally asking the agent to deviate and then rewarding the monitor for detecting this. This would be more difficult to achieve in a doping context since it would require WADA to ask athletes to occasionally take drugs. In addition, Rahman finds that such contracts are not robust to collusion between the monitor and the agent. Finally, Katz (1991) has shown that if the contract between a principal and agent is unobservable, it does not affect the equilibrium outcome if it is common knowledge that there exists a contract which solves agency problems and that the principal and agent have the same preferences over income and effort. This helps to explain why in the baseline case, there is no difference between simultaneous and sequential timing.

All of these papers stress the importance of a higher level organisation, such as the World Anti-Doping Agency, in providing incentives for an organisation to carry out tests. However, this chapter also draws on a paper by Buechel et al. (2014) which considers a sporting scenario and introduces customers as a third player. Previous papers concerning cheating in sport only considered the interaction between athletes and organisers. Buechel et al. use the standard assumption that the ADA prefers to test if the athlete is doped and prefers not to test if the athlete is not doped, reflecting a benevolent ADA. My model allows for the possibility that the ADA prefers not to test if the athlete is doped and may prefer to test if the athlete is clean. This reflects my assumption of a non-benevolent testing agency. In Buechel et al. if customers cannot observe whether or not tests were carried out, then the unique equilibrium is that athletes take drugs and the ADA does not test. This conclusion is echoed in my baseline case. In their model, if instead customers can observe whether or not tests were carried out, there is a no-cheating equilibrium, although this equilibrium is not unique. A unique equilibrium is only found if customers have preferences such that they prefer to withdraw their support if no tests are carried out.

The model presented here extends that of Buechel et al. in several ways. Firstly, Buechel et al. only consider simultaneous timing, while my model suggests sequential timing may be advantageous. Secondly, by specifying general forms for utility functions, it is possible to see their assumptions about preferences in terms of parameter values. Thirdly, their paper only considers the role of customers in deterring cheating and not the potential role of an external testing body. Fourthly, Buechel et al. predominantly analyse consumers who are not fully rational. Finally in my model, the possibility of imperfect tests is accounted for while Buechel et al. only briefly discuss this possibility.

Two other papers which are related are by Eber (2002) and Emrich and Pierdzioch (2013*b*). Eber (2002) uses a Barro-Gordon style framework to show that if anti-doping agencies suffer disutility when athletes are disqualified, they will exert less effort in deterring cheating than is optimal. Eber assumes that athletes choose a level of doping in order to maintain a constant probability of being caught, regardless of the effort exerted by the anti-doping agency. However, the decision of the athlete is not analysed beyond this and it is not explained where this constant probability comes from. Eber also does not allow for the possibility that the anti-doping agency may benefit from the higher quality of performance produced by an athlete who takes drugs. He concludes that the solutions to the credibility problem require either an anti-doping agency which is infinitely doping averse or forcing the anti-doping agency to commit to its announced effort level, by making wages contingent on implementing this level. In my set-up pre-announcement alone is not sufficient and instead, either a higher-level anti-doping agency is required or consumers (who are not considered in Eber's model) must act as an enforcement mechanism.

Meanwhile, Emrich and Pierdzioch (2013*b*) focus on the interaction between two national anti-doping agencies. They find that if a national anti-doping agency cares about the performance of its athletes relative to a rival country, and doping improves performance, the level of doping will be higher than is optimal. Since Emrich and Pierdzioch focus on the interaction between two anti-doping authorities, they do not consider the incentives of athletes regarding whether or not to dope but instead assume that the anti-doping authorities can implement the level of doping which they desire. Consequently, the paper does not consider the role which simultaneous and sequential timing can play in deterring cheating. Furthermore, the only solution which is suggested is to have a monopoly anti-doping agency; the potential roles of consumers or an external testing authority are not considered.

Eber (2002) and Emrich and Pierdzioch (2013*b*) show that there are two reasons why an anti-doping agency may wish to implement a level of doping which is overly high. The first, which is found in Eber (2002), is if the anti-doping agency suffers a loss when an athlete is caught cheating. The second, which is the mechanism in Emrich and Pierdzioch (2013*b*), is if the anti-doping agency cares about the level of performance produced by the athlete. My model incorporates both of these mechanisms in the profit function of the anti-doping agency and then combines this with an analysis of the behaviour of athletes. This allows a more detailed analysis of potential solutions to the problem of non-benevolent testing authorities to be carried out.

One of the key assumptions of this chapter is that consumers care about doping and are not willing to participate if they believe doping is likely to have occurred. Cisyk and Courty (2015) use evidence from Major League Baseball to assess the extent to which the announcement of a doping violation decreases fan attendance. They find that doping decreases consumer demand and that the announcement of a doping offence by a team member costs the team $451,000 in foregone revenue.

3.3 SET-UP

There are three players: an athlete, an anti-doping agency (ADA) and consumers. The analysis is restricted to the case where there is only one athlete. In the game the athlete can decide whether or not to cheat. Cheating improves the athlete's performance but if he is caught he is fined. The ADA decides the probability with which to test the athlete. Testing is costly for the ADA but the ADA may wish to test with a high enough probability that athletes are deterred from cheating. Consumers make a binary decision about whether or not to participate. Consumers benefit from better performances but are only willing to participate if they believe the probability that the athlete cheated is sufficiently low.

There are two types of athlete: type H and type L. An athlete is type H with probability μ. If consumers participate, the athlete and the ADA receive payoffs according to the performance of the athlete. However, if consumers choose not to participate, neither the athlete nor the ADA receive any benefit from the athlete's performance. Therefore the athlete and the ADA only receive the following payoffs if consumers participate. If an athlete of type H does not take drugs, he produces a performance of $\underline{P}_H$. The athlete derives utility a_H and the ADA a payoff of c_H. If type H takes drugs, he produces a performance of $\overline{P}_H$. He derives utility $a_H + b$ and the ADA receives $c_H + d$. If type L does not take drugs, he produces a performance of $\underline{P}_L < \underline{P}_H$ and derives utility a_L while the ADA obtains c_L. If type L takes drugs, he produces a performance of $\overline{P}_L = \underline{P}_H$ and derives utility $a_L + b = a_H$.[1] The ADA receives $c_L + d = c_H$. Therefore it is assumed that a type L who cheats produces the same performance as a type H who does not. This means that if consumers observe this performance level and no additional information, they cannot be certain about the athlete's type which motivates consumers to update their beliefs after observing $\underline{P}_H$.

1 It is assumed that both types of athlete derive the same additional utility from taking drugs. This is not a necessary assumption, but renders the analysis more tractable.

It is assumed that the athlete knows his own type prior to making his decision but consumers and the ADA cannot observe the athlete's type. This means the ADA cannot make its decision about testing probability contingent on type. If the ADA does not commit to a single testing probability, then consumers cannot check the ADA is indeed testing with the probability it announces since consumers cannot observe the athlete's type.

The ADA can choose the probability with which an athlete is tested, t. The selection of testing probability is a pure strategy, not a mixed strategy, for the ADA. This reflects the idea that with multiple athletes or over a long time period, the ADA decides how many athletes to test or how frequently to test the athlete which in the one shot set-up presented here is akin to the ADA choosing the probability with which to test the athlete. Carrying out a test costs the ADA K. It is assumed that $K < c_L$ so that even if the ADA tests the athlete with certainty and the athlete does not cheat, the ADA's participation constraint is still satisfied. The main role of K is to act as a tie-breaker so that for given behaviour by the athlete and consumers, the ADA prefers to choose a lower t rather than a higher one.

If an athlete is tested, false negatives can arise. ρ is the probability of a true positive if the athlete is tested. It is assumed that the probability of a false positive is 0. This allows more concrete results to be derived than would otherwise be the case, but also reflects the fact that the probability of a false positive is likely to be very small. For example, the World Anti-Doping Agency has stated that for the test for human growth hormone, the false positive rate is at least 1 in 10,000 (Macur, 2011). If the athlete tests positive the ADA suffers a lump sum loss of B and the athlete a lump sum loss of F. These losses can either reflect reputational losses or fines. The outside option of both the athlete and the ADA is normalised to zero.

Consumers care about their expected utility which in turn depends on the probability that the athlete cheated to produce the performance they observe. Throughout the analysis it is assumed that all consumers care about doping. The results would not change substantially if a subset of consumers do not care about doping, provided this subset is small. While consumers benefit from better performances, they also dislike doping and are only prepared to participate if they believe the probability that the athlete doped is sufficiently low. Consumers are always able to observe the performance of the athlete and know if the athlete tested positive. In the baseline case, where they observe no additional information about testing, and section 3.5, where they only observe testing probability, they do not know whether the absence of a positive test reflects a negative test or that the athlete was not tested. This is in keeping with the policy of anti-doping agencies which is to only announce positive tests in order to respect the privacy of athletes and ensure that tests remain unpredictable. However, if consumers can only observe positive tests, it may be hard for them to be able to ensure that the ADA is indeed testing with the probability which it announces. Therefore in section 3.6 consumers are also able to observe negative tests.

The expected utility function of consumers is $f(P, \gamma)$ where $P \in \{\underline{P}_L, \underline{P}_H, \overline{P}_H\}$ is the performance of the athlete and γ is the posterior probability that the athlete cheated. The justification for this functional form is given in appendix B.1. Consumers participate if their utility from doing so exceeds their reservation utility which is normalised to 0. It is assumed that $f(P, 1) < 0$ for all performances. This means that consumers never wish to participate if they know with certainty that the athlete cheated. Additionally, $f(P, 0) > 0$ for all P so that if consumers know with certainty that the athlete did not cheat, they will participate for all possible performances. It is also assumed that $f(\underline{P}_H, 1 - \mu) < 0$ so that if consumers cannot update their prior after observing $\underline{P}_H$, they do not wish to participate. Finally, it is assumed that f is monotonically decreasing in γ so that for a given performance, there is a cutoff posterior probability of cheating below which consumers are willing to participate. When the performance is $\underline{P}_H$, this cutoff probability is defined to be $\hat{\gamma}$.

Since the athlete can only test positive if he cheated, consumers never wish to participate after observing a positive test. If consumers observe $\underline{P}_L$, they know with certainty that this performance was generated by a type L athlete who did not cheat and therefore participate. If consumers observe performance $\overline{P}_H$, they know with certainty that this performance was generated by a type H athlete who cheated and therefore prefer not to participate. However, consumers cannot infer from performance $\overline{P}_L = \underline{P}_H$ alone whether or not the athlete cheated. Therefore after observing this performance they engage in Bayesian updating and use the resulting posterior probability to determine whether or not to participate.

Athletes and the ADA are expected utility maximisers. If type L does not cheat, consumers definitely participate and there is no possibility of the athlete testing positive. His utility is therefore a_L. If type L cheats, his utility depends on both the testing probability and whether or not consumers decide to participate. If consumers do not participate his expected utility is $-\rho t F$. Therefore if type L expects consumers not to participate if he cheats, he prefers not to cheat for all values of t. This is because when consumers do not participate the athlete does not get the benefit from his performance but can still test positive if he cheats. If type L cheats and consumers participate provided he does not test positive, type L obtains:

$$(1-\rho t)(a_L+b)-\rho t F$$

Type L tests positive with probability ρt. Provided he does not test positive he obtains a_L+b, but if he tests positive he is also fined F. Therefore if type L expects consumers to participate if he cheats provided he does not test positive, he prefers not to cheat if:

$$t \geq \frac{b}{\rho(a_L+b+F)} \equiv t_1$$

Type H does not cheat regardless of the behaviour of either consumers or the ADA. When type H cheats, he produces a performance of $\overline{P}_H$ and consumers recognise this performance must be the result of cheating and therefore do not participate. Consequently if type H cheats his expected utility is weakly negative since if $t>0$, there is the possibility that he will test positive. If consumers participate when he does not cheat, his utility is strictly positive. If they do not participate it is equal to zero. Therefore type H does not cheat regardless of the t chosen by the ADA and whether or not consumers participate when he does not cheat.[2]

2 If consumers do not participate when type H does not cheat and $t=0$, type H obtains zero utility regardless of whether or not he cheats. Throughout the analysis it will be assumed that in this case type H does not cheat. This is a natural assumption and could for example reflect unmodelled health costs of doping.

In summary, if $t \geq t_1$, neither type cheats regardless of the behaviour of consumers. If $t < t_1$, type H never cheats. Type L cheats if consumers participate and does not cheat if they do not.

The solution concept used throughout this chapter is perfect Bayesian equilibrium. However, Fudenberg and Tirole (1991*b*) show that when each player has at most two possible types, the set of perfect Bayesian equilibria and the set of sequential equilibria coincide.

3.4 BASELINE CASE

In the baseline case, consumers can only update their belief that the athlete cheated based on the performance they observe and the lack or presence of a positive test. They cannot observe the testing probability. The interesting case is when they observe performance $\overline{P}_L = \underline{P}_H$ since this is the only case in which there is doubt about whether or not the athlete cheated.

Proposition 3.1. *In the baseline case, regardless of whether timing is simultaneous or sequential, there is no perfect Bayesian equilibrium in which neither type of athlete cheats.*

Proof. See appendix B.2. □

Attention is restricted to the existence of an equilibrium in which type L cheats with zero probability. The actions taken on the equilibrium path to generate a no-cheating equilibrium must be as follows:

- ADA: $t = t_1$
- Type H athlete: Don't cheat
- Type L athlete: Don't cheat
- Consumers: Participate

In any no-cheating equilibrium, consumers must participate with certainty after observing $\underline{P}_H$ since they then know with certainty that the performance was produced by an athlete of type H who did not cheat. If consumers participate after observing $\underline{P}_H$, type L will cheat unless $t \geq t_1$. Therefore the existence of a no-cheating PBE in the baseline case depends on whether or not the ADA finds it optimal to set $t = t_1$.

When timing is simultaneous, neither consumers nor the athlete observe testing probability prior to making their decision. The ADA then has no incentive to test with strictly positive probability. The ADA takes the decision of the athlete and the consumer as given and therefore prefers to deviate to $t = 0$. By doing so it avoids having to carry out costly tests while its other expected payoffs are unchanged. However, we cannot have a no-cheating equilibrium in which consumers participate and $t = 0$ since type L then prefers to cheat.

When timing is sequential, the same set of strategies is required on the equilibrium path for a no-cheating equilibrium as when timing is simultaneous. The difference is that now, the ADA recognises that type L will only wish not to cheat with certainty if $t \geq t_1$. However, when consumers cannot make their strategy contingent on the observed testing probability, the ADA always has an incentive to deviate and set $t = 0$. If the ADA expects consumers to participate it benefits from deviating to $t = 0$ since testing is costly, the ADA benefits when type L cheats and produces a higher performance and when $t = 0$ there is no risk of the athlete testing positive. Therefore an equilibrium with $t = t_1$ does not exist when timing is sequential, but as was seen above, $t = t_1$ is a requirement for any no-cheating equilibrium.

Consequently, full rationality of consumers is not enough to ensure the existence of an equilibrium in which neither type of athlete cheats if consumers cannot observe the probability with which the athlete was tested. The difficulty is that even if the ADA wanted to implement a no-cheating equilibrium, without observability of t, it cannot credibly commit itself to setting t_1. Consequently, in the next section I examine the case where consumers can observe testing probability.

3.5 CONSUMERS OBSERVE TESTING PROBABILITY

Full rationality of consumers is not enough to ensure the existence of an equilibrium in which no type of athlete cheats if consumers cannot observe the probability with which the athlete was tested. Therefore this section extends the analysis to allow for the possibility that consumers can observe the testing probability prior to making their decision regarding whether or not to participate. However, consumers still cannot distinguish between a negative test and no test having been carried out

Once again, consumers participate provided the probability that the athlete cheated is sufficiently low. The utility function of consumers is again $f(P, \gamma)$ where γ is the posterior probability that the athlete cheated. The only difference compared to the baseline case is that consumers' posterior probability can now be contingent on the probability of testing which they observe. As above, $\hat{\gamma}$ is defined so that $f(\underline{P}_H, \hat{\gamma}) = 0$.

As in the baseline case, if the athlete tests positive or produces performance $\overline{P}_H$, consumers can infer with certainty that the athlete cheated and therefore do not participate. If consumers observe $\underline{P}_L$, they know with certainty that this performance was generated by a type L athlete who did not cheat and therefore participate. After observing $\underline{P}_H$, t and that the athlete did not test positive, consumers use Bayes' rule to update their belief that the athlete cheated. This belief coincides with their belief that the athlete is type L since type L cheats to produce $\underline{P}_H$. Since in the baseline case there is no equilibrium in which neither type cheats with certainty, attention is focused on the existence of such an equilibrium.

3.5.1 *Simultaneous Timing*

The timing of the game is as follows:

1. The athlete decides whether or not to cheat. Simultaneously the ADA decides the probability with which it will test the athlete. After these decisions have been made, consumers observe the athlete's performance and the t chosen by the organisation.

2. With probability t, the athlete is tested. If the athlete tests positive, consumers observe this, but in the absence of a positive test, they cannot observe whether or not the athlete was tested.

3. Consumers choose whether or not to participate, based on the information available to them.

Consumers use Bayes' rule to update their beliefs about the athlete's type at the end of each period. At the start of the first period, consumers have no new information so believe the athlete is type L with probability $1 - \mu$. At the start of the second period, consumers update their belief based on the performance they observe. $h^{\mathrm{t}-1}$ is the history at the start of period t and a^{t} is the vector of actions taken in period t. If consumers observe performance $\underline{P}_H$, at the start of period 2 their belief is:

$$\Pr(\text{Type L} \mid h^0,\, a^1) \quad = \quad \frac{(1-\mu)\Pr(\underline{P}_H \mid \text{Type L},\, h^0)}{(1-\mu)\Pr(\underline{P}_H \mid \text{Type L},\, h^0) + \mu}$$

since $\Pr(\underline{P}_H \mid \text{Type H},\, h^0) = 1$. Fudenberg and Tirole (1991*b*) state that in a Perfect Bayesian Equilibrium, 'the beliefs about a player at the beginning of period (t+1) depend only on the history up to date t and on that player's date-t action, but not on the other players' date-t actions'. Therefore, the ADA's choice of t does not change consumers' belief that the athlete is type L at the start of the second period. Furthermore, in a perfect Bayesian equilibrium, Bayes' rule is used to update beliefs about player j from period t to t+1 if the history at the start of period t had positive probability and some player $k \neq j$ chooses a probability-0 action at date t (Fudenberg and Tirole, 1991*a*). In the set-up here, this means that even if the ADA chooses a probability-0 action in period 1, consumers still use Bayes' rule to update their beliefs about the athlete's type.

At the start of the third period, consumers can update again based on whether or not they observe that the athlete tested positive. If they observe that the athlete did not test positive:

$$\begin{aligned}\gamma &= \Pr(\text{Type L} \mid h^1, \text{ no positive}) \\ &= \frac{\Pr(\text{Type L} \mid h^0,\ a^1)\Pr(\text{no positive} \mid \text{Type L},\ h^1)}{\Pr(\text{Type L} \mid h^0,\ a^1)\Pr(\text{no positive} \mid \text{Type L},\ h^1) + \mu} \\ &= \frac{(1-\mu)(1-\rho t)\Pr(\underline{P}_H \mid \text{Type L},\ h^0)}{(1-\mu)(1-\rho t)\Pr(\underline{P}_H \mid \text{Type L},\ h^0) + \mu}\end{aligned}$$

Proposition 3.2. *When timing is simultaneous, a no-cheating perfect Bayesian equilibrium does not exist.*

We are interested in the existence of an equilibrium in which both types of athlete have a pure strategy of not cheating. The strategies which would be required for a no-cheating equilibrium are as follows:

- ADA: $t = t_1$
- Type H athlete: Don't Cheat
- Type L athlete: Don't Cheat
- Consumers: Participate if $\underline{P}_L$ or $\underline{P}_H$ and athlete does not test positive regardless of testing probability; don't participate if $\overline{P}_H$ or athlete tests positive

When timing is simultaneous, the athletes' strategies cannot be contingent on the t chosen by the ADA since athletes do not observe t prior to making their decision. The difference compared to the baseline case is that consumers' strategy can now be contingent on the t chosen by the ADA. In any no-cheating equilibrium, consumers must participate on the equilibrium path. A no-cheating equilibrium then exists provided it is optimal for the ADA to set $t = t_1$. This in turn depends on how consumers behave after observing the off-equilibrium path information set $\{t < t_1, \underline{P}_H, \text{no positive test}\}$. If consumers do not participate after observing this information set, it will not be optimal for the ADA to deviate and a no-cheating equilibrium exists. If instead consumers participate with certainty when $t < t_1$, the ADA always prefers to set $t = 0$ since testing is costly.

If the ADA chooses the probability-0 action $t < t_1$ in period 1, consumers recognise that the athlete did not observe this prior to making his decision and therefore follows his strategy of not cheating. Then $\Pr(\underline{P}_H \mid \text{Type L}, h^0) = 0$ and consumers attach zero probability to the athlete being type L after observing $\{t < t_1, \underline{P}_H, \text{no positive test}\}$. γ then equals zero and consumers participate. However, if consumers participate after observing $\{t < t_1, \underline{P}_H, \text{no positive test}\}$, the ADA always wishes to deviate and set $t = 0$ since then the ADA need not undertake costly testing and its payoff from both types of athlete is unchanged. This then breaks the no-cheating equilibrium.

3.5.2 *Sequential Timing*

The timing of this game is as follows:

1. The ADA informs the athlete and the consumer of the probability with which the athlete will be tested.
2. The athlete decides whether or not to cheat. Consumers observe the athlete's performance.
3. With probability t, the athlete is tested. If the athlete tests positive, consumers observe this, but in the absence of a positive test, they cannot observe whether or not the athlete was tested.
4. Consumers choose whether or not to participate, based on the information available to them.

At the start of period 2, consumers still believe the athlete is type L with probability $1 - \mu$. This is because the ADA does not know the athlete's type and so its choice of t conveys no new information about this. However, after observing the athlete's performance, consumers can update their beliefs.[3] If consumers observe $\underline{P}_H$, at the start of period 3 they believe:

$$\begin{aligned}\Pr(\text{Type L} \mid h^1,\, a^2) &= \frac{\Pr(\text{Type L} \mid h^1)\Pr(\underline{P}_H \mid \text{Type L},\, h^1)}{\Pr(\text{Type L} \mid h^1)\Pr(\underline{P}_H \mid \text{Type L},\, h^1) + \mu} \\ &= \frac{(1-\mu)\Pr(\underline{P}_H \mid \text{Type L},\, h^1)}{(1-\mu)\Pr(\underline{P}_H \mid \text{Type L},\, h^1) + \mu}\end{aligned}$$

After observing that the athlete did not test positive, consumers can further update their beliefs and:

$$\gamma = \Pr(\text{Type L} \mid h^2,\, \text{no positive}) = \frac{(1-\mu)(1-\rho t)\Pr(\underline{P}_H \mid \text{Type L},\, h^1)}{(1-\mu)(1-\rho t)\Pr(\underline{P}_H \mid \text{Type L},\, h^1) + \mu}$$

The key difference between this belief and that held by consumers when timing is simultaneous is that the probability that a type L athlete produces performance $\underline{P}_H$ is now also contingent on the t chosen by the organisation and consumers recognise this.

Proposition 3.3. *When timing is sequential, a unique no-cheating perfect Bayesian equilibrium exists provided condition 3.1 is satisfied for all* $\tilde{t} \in \left[0, min\{\hat{t}, t_1 - \epsilon\}\right]$:

$$\begin{aligned}(1-\mu)\kappa(\tilde{t})\Big[c_L + \rho\tilde{t}B - \sigma(\tilde{t})(1-\rho\tilde{t})(c_L + d)\Big] \\ +\mu\Big(1-\sigma(\tilde{t})\Big)(c_L + d) - (t_1 - \tilde{t})K \geq 0 \qquad (3.1)\end{aligned}$$

Proof. See appendix B.3. □

$\kappa(\tilde{t})$ is the probability with which type L cheats if $t < t_1$.

$$\kappa(\tilde{t}) = \begin{cases} \frac{\mu\hat{\gamma}}{(1-\hat{\gamma})(1-\mu)(1-\rho t)} & \text{if } \tilde{t} < min\{\hat{t}, t_1\} \\ 1 & \text{if } \hat{t} \leq \tilde{t} < t_1 \end{cases}$$

3 In a perfect Bayesian equilibrium, Bayes' rule is used to update beliefs from period t to t+1 even if the history at the start of period t has probability 0. In this set-up therefore, even if the ADA chooses a probability-0 action in period 1, consumers still use Bayes' rule to update their belief about the athlete's type based on his action in period 2 (Fudenberg and Tirole, 1991*a*).

$\sigma(\tilde{t})$ is the probability with which consumers participate after observing the information set $\{t < t_1, \underline{P}_H$, no positive test$\}$.

$$\sigma(\tilde{t}) = \begin{cases} \frac{a_L + \rho t F}{(1-\rho t)(a_L + b)} & \text{if } \tilde{t} < min\{\hat{t}, t_1\} \\ 1 & \text{if } \hat{t} \leq \tilde{t} < t_1 \end{cases}$$

$\kappa(\tilde{t})$ and $\sigma(\tilde{t})$ are derived in the appendix. After observing $\{t < t_1, \underline{P}_H$, no positive test$\}$, consumers are willing to participate if γ, the probability with which they believe the athlete cheated, is less than $\hat{\gamma}$ where:

$$\gamma = \frac{(1-\rho t)(1-\mu)\kappa(t)}{(1-\rho t)(1-\mu)\kappa(t) + \mu}$$

It is possible that there exists a t sufficiently high that consumers are willing to participate with certainty after observing this t even if type L cheats with certainty $\big(\kappa(t) = 1\big)$. Define this t to be $\hat{t}$ where $\hat{t}$ is defined such that:

$$\hat{\gamma} = \frac{(1-\rho \hat{t})(1-\mu)}{(1-\rho \hat{t})(1-\mu) + \mu}$$

Intuitively, if t is high, there is a high probability that a type L athlete will test positive if he cheats. Therefore the absence of a positive test acts as a sufficiently good signal that the athlete is not type L that consumers are willing to participate. If $\hat{t} < t_1$, there exists a range of $t \in [\hat{t}, t_1)$ where after observing this t type L cheats with certainty and consumers participate with certainty.

The equilibrium strategies are as follows. Actions taken on the equilibrium path are highlighted in bold:

- ADA: $\boldsymbol{t = t_1}$
- Type H athlete: **Don't Cheat**
- Type L athlete: Cheat with probability $\kappa(t)$ if $t \in \Big[0, min\{\hat{t}, t_1\}\Big)$; cheat if $t \in [\hat{t}, t_1)$; **don't cheat if $\boldsymbol{t \geq t_1}$**

- Consumers: **Participate if $\underline{P}_L$ regardless of testing**; participate with probability $\sigma(t)$ if $\underline{P}_H$ and $t < min\{\hat{t}, t_1\}$; **participate if $\underline{P}_H$ and $t \geq min\{\hat{t}, t_1\}$**; don't participate if $\overline{P}_H$ or athlete tests positive regardless of testing probability.

- Consumers' beliefs: **$\gamma = 0$ if $\underline{P}_L$, $\gamma = 0$ if $\underline{P}_H$ and $t \geq t_1$**. $\gamma = \frac{(1-\mu)(1-\rho t)}{(1-\mu)(1-\rho t)+\mu}$ if $\underline{P}_H$ and $\hat{t} \leq t < t_1$. $\gamma = \hat{\gamma}$ if $t \in \left[0, min\{\hat{t}, t_1\}\right)$ and $\gamma = 1$ if $\overline{P}_H$ or the athlete tests positive.

As was discussed above, in any potential no-cheating equilibrium it must be the case that on the equilibrium path, neither type of athlete wishes to cheat, consumers participate and the ADA sets $t = t_1$. When timing was simultaneous such an equilibrium could not exist because given the strategies of the other players, the ADA always had an incentive to deviate and set $t = 0$. However, when timing is sequential, type L's strategy is contingent on the testing probability announced by the ADA. Both the ADA and consumers therefore recognise that if the ADA deviates to $t = 0$, type L will cheat with positive probability. Consumers then do not wish to participate with certainty after observing $\{t = 0$, $\underline{P}_H$, no positive test$\}$.

If $t < min\{t_1, \hat{t}\}$, beliefs and strategies cannot be consistent if attention is restricted to pure strategies. If type L expects consumers to participate with certainty if he cheats, he will cheat but then consumers do not wish to participate. If type L expects consumers not to participate with certainty if he cheats he will not cheat, but then consumers wish to participate after observing $\underline{P}_H$. Consistent beliefs and strategies can only be found if consumers are able to mix between participating and not while type L mixes between cheating and not. $\kappa(t)$ is the probability of type L cheating which makes consumers indifferent between participating and not after observing $\underline{P}_H$, t and that the athlete did not test positive when $t \in \left[0, min\{\hat{t}, t_1\}\right)$. $\sigma(t)$ is the probability of consumers participating after this information set which makes type L indifferent between cheating and not. These probabilities are derived in the appendix.

A no-cheating equilibrium exists provided the ADA prefers to set $t = t_1$ rather than $t < t_1$. The ADA's payoff from setting $t = t_1$ and having neither type of athlete cheat with certainty is greater than its payoff from setting $\tilde{t} < t_1$ and having type L cheat with positive probability if:

$$\begin{aligned}(1-\mu)\kappa(\tilde{t})\left[c_L + \rho\tilde{t}B - \sigma(\tilde{t})(1-\rho\tilde{t})(c_L+d)\right] \\ +\mu\Big(1-\sigma(\tilde{t})\Big)(c_L+d) - (t_1-\tilde{t})K \quad \geq \quad 0\end{aligned}$$

for all $\tilde{t} \in \Big[0, min\{t_1 - \epsilon, \hat{t}\}\Big]$. The first term refers to the potential gain from setting t_1 rather than $\tilde{t}$ if the athlete is type L. This term may be positive or negative depending on the exact parameter values. The second term refers to the gain from setting t_1 rather than $\tilde{t}$ if the athlete is type H. This term is unambiguously positive.

As K, the cost of testing, tends to zero, the only benefit for the ADA from setting $t < t_1$ occurs if the athlete turns out to be type L and the first term of condition 3.1 is negative. If this is the case, the ADA gains d from setting $t < t_1$ rather than $t = t_1$ if the athlete cheats, does not test positive and consumers participate. However, if the athlete cheats and consumers do not participate or the athlete tests positive, the ADA loses c_L. Additionally, if the athlete cheats and tests positive the ADA loses B. If type L does not cheat, there is no difference in the ADA's payoff when $t < t_1$ and $t = t_1$. If the athlete is type H, the ADA is guaranteed $c_L + d$ if it sets t_1. However, if it sets $t < t_1$ there is a risk that consumers will not participate after observing $\underline{P}_H$ and the ADA then loses $c_L + d$.

The left hand side of the above condition is strictly decreasing in $\sigma(\tilde{t})$. If consumers are less likely to participate after observing $\underline{P}_H$ and $t < t_1$, the ADA finds it less appealing to set a low value of t. In the appendix it is shown that if $0 < t < min\{\hat{t}, t_1\}$, $\sigma(\tilde{t})$ is increasing in F and decreasing in b. Therefore when F increases, the left hand side of condition 3.1 decreases and the ADA is more likely to wish to set $t < t_1$. At first it appears unintuitive that as the punishment for an athlete caught cheating increases, a no-cheating equilibrium becomes less likely. However, the reason for this is that as F increases, consumers participate with a higher probability after observing $\underline{P}_H$ and $0 < t < min\{\hat{t}, t_1\}$ and this makes it more appealing for the ADA not to implement the no-cheating equilibrium. The same intuition applies to why an increase in b makes it more likely that a no-cheating equilibrium exists. For $\tilde{t} > 0$, condition 3.1 is increasing in B. Intuitively, if the fine the ADA must pay when the athlete tests positive increases, the ADA prefers to implement a no-cheating equilibrium for a wider range of parameter values.

For all $\tilde{t} > 0$, a no-cheating equilibrium is therefore guaranteed provided B is sufficiently high. The ADA then does not wish to set $t \in (0, t_1)$ since there is a risk type L tests positive and the ADA must pay a large fine. However, when $\tilde{t} = 0$, there is no risk of type L testing positive and therefore a large B alone cannot guarantee the existence of a no-cheating equilibrium. A necessary condition for the ADA to prefer to set $t = t_1$ rather than $t = 0$ as $K \to 0$ is that:

$$\hat{\gamma} \leq \frac{1 - \sigma(0)}{d/(c_L + d)} = \frac{b/(a_L + b)}{d/(c_L + d)} \equiv \bar{\gamma}$$

This shows that a no-cheating equilibrium with $t = t_1$ is preferred to a cheating equilibrium with $t = 0$ provided consumers are only prepared to tolerate a low probability of cheating. As $\sigma(0)$, the probability with which consumers participate when $t = 0$, increases, $\bar{\gamma}$ decreases meaning $\hat{\gamma} < \bar{\gamma}$ for fewer values of $\hat{\gamma}$. The intuition for this is the same as above. When d increases, so that the ADA derives a greater benefit when type L cheats and consumers participate, $\bar{\gamma}$ decreases meaning a no-cheating equilibrium occurs for fewer values of $\hat{\gamma}$. An increase in c_L, the ADA's benefit when type L does not cheat, has the opposite effect.

When the parameter values are such that a no-cheating equilibrium exists, this equilibrium is always unique. This is because Bayes' rule is used by consumers to update their belief about the athlete even when the ADA takes a probability-0 action. Consumers have a unique consistent belief after each information set and their optimal action after each information set is therefore also unique. Consequently, there is a unique testing probability which maximises the ADA's profit, taking into account the strategies of the other players.

3.5.2.1 *Proportional Payoffs*

Further analysis can be carried out if we make the simplifying assumption that $c_L = \eta a_L$, $d = \eta b$ and $B = \eta F$ where $\eta > 0$. This means that the ADA's payoffs are all proportional to the athlete's. This simplifies the analysis since it is then possible to compare the athlete's payoffs, which enter into $\sigma(t)$, with the ADA's. The profit the ADA then derives from deviating to $t < t_1$ depends on the value of $min\{\hat{t}, t_1\}$. If $min\{\hat{t}, t_1\} = t_1$, as $K \to 0$, the ADA's profit from deviating to $\tilde{t} < t_1$ is:

$$(1-\mu)c_L + \mu\sigma(\tilde{t})(c_L + d)$$

which is increasing in t regardless of B. Therefore if the ADA chooses to deviate, it does so by setting $t_1 - \epsilon$. The profit from the ADA deviating is then always less than its profit from setting $t = t_1$. The ADA obtains the same profit from type L regardless of whether it deviates or not. This is because when payoffs are proportional, the value of $\sigma(t)$ which makes type L indifferent between cheating and not also makes the ADA indifferent between whether or not type L cheats. However, the ADA derives strictly less profit from deviating if the athlete is type H. If the ADA does not deviate, consumers always participate and the ADA is guaranteed $c_L + d$. However, if the ADA deviates, there is a chance that consumers do not participate after observing $\underline{P}_H$ and the ADA then receives no payoff from a type H athlete. Therefore, if $min\{\hat{t}, t_1\} = t_1$, the ADA will never wish to deviate. This is proved formally in appendix B.3.1.

If instead $min\{\hat{t}, t_1\} = \hat{t}$, as $K \to 0$, the ADA's profit from setting $\hat{t}$ is:

$$(1-\mu)\Big[(1-\rho\hat{t})(c_L + d) - \rho\hat{t}B\Big] + \mu(c_L + d)$$

In the appendix it is shown that given this profit function, if the ADA wishes to deviate it will always do so by setting $\hat{t}$. If the athlete is type H, the ADA is then guaranteed $c_L + d$ regardless of whether or not it deviates. If the athlete is type L, the ADA obtains $(1 - \rho\hat{t})(c_L + d) - \rho\hat{t}B$ if it deviates and c_L if it does not. When payoffs are proportional, the same value of t which makes the athlete indifferent between cheating and not if consumers participate, t_1, also makes the ADA indifferent between deviating and not. Therefore, if consumers are prepared to participate with certainty when the ADA sets $\hat{t} < t_1$, the ADA will always wish to deviate. This is proved formally in appendix B.3.1. However, for B

sufficiently large, it is still the case that a no-cheating equilibrium exists. This is because as B increases, t_1 decreases reducing the range of parameter values for which $\hat{t} < t_1$.

3.6 CONSUMERS OBSERVE NEGATIVE TESTS

So far it has been assumed that observability for consumers means they are able to observe the probability with which the athlete was tested and if the outcome of the test was positive. However, in the previous sections, in the absence of a positive test, consumers could not distinguish between a negative test and the athlete having not been tested. In this section, consumers are able to distinguish between a negative test and no test having been carried out. It is important to check the results above are robust to this extension of the information sets available to consumers since if consumers cannot observe negative tests, it may not be possible for them to check the ADA is testing with the probability it announces. In addition, it was seen above that when payoffs are proportional and consumers cannot observe negative tests, a no-cheating equilibrium does not exist if $min\{t_1, \hat{t}\} = \hat{t}$. However, when consumers can observe negative tests, a no-cheating equilibrium always exists when payoffs are proportional.

If consumers observe that performance is $\underline{P}_H$ and the athlete was tested with probability t, they are able to distinguish between three information sets:

1. $\{t, \underline{P}_H,$ positive test$\}$
2. $\{t, \underline{P}_H,$ negative test$\}$
3. $\{t, \underline{P}_H,$ no test$\}$

This is a finer information structure than was previously available to consumers since previously consumers could not distinguish between cases two and three. As was previously the case, an athlete who does not cheat cannot test positive. Therefore after observing the first information set, consumers know with certainty the athlete was a type L who cheated and therefore do not participate. Consumers may or may not be willing to participate after observing the second and third information sets.

3.6.1 *Simultaneous Timing*

The timing of this game is as follows:

1. The athlete decides whether or not to cheat. Simultaneously the ADA decides the probability with which it will test the athlete. After these decisions have been made, consumers observe the athlete's performance and the t chosen by the organisation.
2. With probability t, the athlete is tested. Consumers observe whether or not the athlete was tested, and if he was, whether the test was positive or negative.
3. Consumers choose whether or not to participate, based on the information available to them.

After observing $\{t,\ \underline{P}_H,\ \text{negative test}\}$, consumers therefore believe the athlete is type L with probability:

$$\gamma = \frac{(1-\mu)(1-\rho)\Pr(\underline{P}_H \mid \text{Type L},\ h^0)}{(1-\mu)(1-\rho)\Pr(\underline{P}_H \mid \text{Type L},\ h^0)+\mu}$$

After observing $\{t,\ \underline{P}_H,\ \text{no test}\}$, consumers believe the athlete is type L with probability:

$$\gamma = \frac{(1-\mu)\Pr(\underline{P}_H \mid \text{Type L},\ h^0)}{(1-\mu)\Pr(\underline{P}_H \mid \text{Type L},\ h^0)+\mu}$$

since consumers obtain no new information after period 2 if the athlete is not tested.

Proposition 3.4. *When timing is simultaneous, a no-cheating perfect Bayesian equilibrium does not exist.*

Proof. The proof of this is exactly the same as in the case where consumers cannot observe negative tests. □

For a no-cheating equilibrium to exist, the following actions are required on the equilibrium path:

- ADA: $t \geq t_1$
- Type H athlete: Don't Cheat
- Type L athlete: Don't Cheat
- Consumers: Participate if $\underline{P}_L$; participate if $\underline{P}_H$ and athlete tests negative or is not tested

Once again, in any no-cheating equilibrium, on the equilibrium path consumers must participate after observing $\underline{P}_H$ since when type L does not cheat, this performance can only be produced by a type H who does not cheat. Therefore, consumers must participate after observing $\underline{P}_H$, regardless of whether the athlete tested negative or was not tested. When consumers participate after observing $\underline{P}_H$ provided the athlete does not test positive, type L will cheat unless $t \geq t_1$.

As in section 3.5.1 a no-cheating equilibrium therefore exists if it is optimal for the ADA to set $t = t_1$. This in turn depends on whether or not consumers participate after observing the information set $\{t < t_1, \underline{P}_H, \text{negative test}\}$ or $\{t < t_1, \underline{P}_H, \text{no test}\}$. However, consumers recognise that if the ADA deviates and sets $t < t_1$, this is not observed by the athlete at the time of making his decision. They therefore believe the athlete followed his strategy of not cheating even if the ADA sets $t < t_1$ and are willing to participate regardless of whether the athlete tests negative or was not tested. The ADA then wishes to deviate to $t = 0$ and this breaks the no-cheating equilibrium.

3.6.2 *Sequential Timing*

The timing of this game is as follows:

1. The ADA informs the athlete and the consumer of the probability with which the athlete will be tested.
2. The athlete decides whether or not to cheat. Consumers observe the athlete's performance.

3. With probability t, the athlete is tested. Consumers observe whether or not the athlete was tested, and if he was, whether the test was positive or negative.

4. Consumers choose whether or not to participate, based on the information available to them.

After observing the information set $\{t,\ \underline{P}_H,\ \text{negative test}\}$, consumers believe the athlete is type L with probability:

$$\gamma \;=\; \frac{(1-\mu)(1-\rho)\Pr(\underline{P}_H \mid \text{Type L},\ h^1)}{(1-\mu)(1-\rho)\Pr(\underline{P}_H \mid \text{Type L},\ h^1)+\mu}$$

After observing $\{t,\ \underline{P}_H,\ \text{no test}\}$, consumers believe the athlete is type L with probability:

$$\gamma \;=\; \frac{(1-\mu)\Pr(\underline{P}_H \mid \text{Type L},\ h^1)}{(1-\mu)\Pr(\underline{P}_H \mid \text{Type L},\ h^1)+\mu}$$

Proposition 3.5. *When timing is sequential, a unique no-cheating perfect Bayesian equilibrium exists provided condition 3.2 is satisfied for all $\tilde{t} \in [0, t_1)$:*

$$\begin{aligned}(1-\mu)\kappa_{NT}\Big[c_L+\rho\tilde{t}B-\Big(\sigma_{NT}(\tilde{t})(1-\tilde{t})+(1-\rho)\tilde{t}\Big)(c_L+d)\Big]&\\ +\mu(1-\tilde{t})\Big(1-\sigma_{NT}(\tilde{t})\Big)(c_L+d)-(t_1-\tilde{t})K \;&\geq\; 0 \qquad (3.2)\end{aligned}$$

Proof. A sketch proof is provided below. A full proof is provided in the appendix and $\sigma_{NT}(\tilde{t})$ and κ_{NT} are derived in the appendix. □

For a no-cheating equilibrium, the following actions must be taken on the equilibrium path:

- ADA: $t \geq t_1$
- Type H athlete: Don't Cheat
- Type L athlete: Don't Cheat

- Consumers: Participate if $\underline{P}_L$; participate after $\{t_1, \underline{P}_H$, negative test$\}$ and $\{t_1, \underline{P}_H$, no test$\}$

If consumers believe that neither type of athlete cheated given the t set by the ADA, they must participate regardless of whether the athlete was not tested or tests negative. Therefore, in any no cheating equilibrium it must be the case that consumers participate after observing $\{t_1,\ \underline{P}_H,\ \text{negative test}\}$ or $\{t_1,\ \underline{P}_H,\ \text{no test}\}$. If consumers participate provided the athlete does not test positive, type L will cheat unless $t \geq t_1$. Consequently, a no-cheating equilibrium exists if it is optimal for the ADA to set $t \geq t_1$.

In considering any possible deviation, the ADA must now consider how consumers will behave after observing both $\{t,\ \underline{P}_H,\ \text{negative test}\}$ and $\{t,\ \underline{P}_H,\ \text{no test}\}$. Previously the ADA did not need to consider these information sets separately since consumers could not distinguish between them.

Consumers will never wish to participate with certainty after both information sets if $t < t_1$ since then type L cheats with certainty. Consumers then cannot update their prior after observing $\{t < t_1,\ \underline{P}_H,\ \text{no test}\}$ and therefore do not wish to participate with certainty. If consumers do not participate after either information set, type L does not cheat and consumers then prefer to participate.

After observing the information set $\{t < t_1,\ \underline{P}_H,\ \text{negative test}\}$, consumers believe the athlete is type L with probability $\gamma_1(t)$ where:

$$\gamma_1(t) = \frac{(1-\mu)(1-\rho)\kappa(t)}{(1-\mu)(1-\rho)\kappa(t) + \mu}$$

and $\kappa(t)$ is the probability with which type L cheats.

After observing $\{t < t_1,\ \underline{P}_H,\ \text{no test}\}$, consumers believe the athlete is type L with probability $\gamma_2(t)$ where:

$$\gamma_2(t) = \frac{(1-\mu)\kappa(t)}{(1-\mu)\kappa(t) + \mu}$$

Consumers participate with strictly positive probability if $\gamma \leq \hat{\gamma}$. The athlete cannot condition his strategy on whether or not he is tested; only on the t. Therefore $\kappa(t)$

is the same in $\gamma_1(t)$ and $\gamma_2(t)$. Since $\gamma_1(t) < \gamma_2(t)$ when $\kappa(t) > 0$, if consumers participate with strictly positive probability after observing the information set $\{t < t_1,\ \underline{P}_H,$ no test$\}$, they must participate with certainty after observing the information set $\{t < t_1,\ \underline{P}_H,$ negative test$\}$.

After these possibilities have been ruled out, this leaves two possible sets of behaviour by consumers if the ADA deviates. These are:

1. Participate with positive probability after observing the information set $\{t < t_1$, $\underline{P}_H$, negative test$\}$; don't participate after observing $\{t < t_1, \underline{P}_H$, no test$\}$.

2. Participate with certainty after observing $\{t < t_1$, $\underline{P}_H$, negative test$\}$; mix after observing $\{t < t_1, \underline{P}_H$, no test$\}$.

For either of these strategies to be consistent, type L must cheat with positive probability if $t < t_1$ since if he does not, consumers wish to participate with certainty after both information sets. If consumers participate with positive probability after observing a negative test but do not participate after no test, it is shown in the appendix that type L does not wish to cheat with positive probability unless $t_1 > 1$. However, when this is the case the ADA cannot implement no-cheating even if it wished to. Therefore, whenever it is possible for the ADA to set t_1, it will never be the case that if the ADA deviates consumers will participate with positive probability after observing $\{t,\ \underline{P}_H,$ negative test$\}$ but not participate after observing $\{t,\ \underline{P}_H,$ no test$\}$. Intuitively, if type L is willing to cheat with positive probability when consumers only participate after observing a negative test, there is no value of t which can deter cheating if consumers participate with certainty provided they do not observe a positive test.

Therefore, if the ADA deviates and does not set $t = t_1$, there is only one set of consistent actions for consumers to take. This is to participate with certainty after observing $\{t,\ \underline{P}_H,$ negative test$\}$ and to mix after observing $\{t,\ \underline{P}_H,$ no test$\}$. Consumers are prepared to do this provided $\kappa(t)$ is such that $\gamma_2 = \hat{\gamma}$. This κ is defined to be κ_{NT}. In turn, given this behaviour by consumers, type L is prepared to mix between cheating and not provided consumers mix with probability $\sigma_{NT}(t)$. In the appendix it is shown that $\sigma_{NT}(t) < 1$ provided $t < t_1$.

Therefore if the ADA deviates and sets $t < t_1$, consumers participate after a negative test, mix after no test and type L mixes between cheating and not. There is no possible profitable deviation involving $t > t_1$ for the ADA. The ADA prefers not to deviate if:

$$
\begin{aligned}
(1-\mu)\kappa_{NT}\left[c_L + \rho\tilde{t}B - \left(\sigma_{NT}(\tilde{t})(1-\tilde{t}) + (1-\rho)\tilde{t}\right)(c_L + d)\right] & \\
+\mu(1-\tilde{t})\left(1-\sigma_{NT}(\tilde{t})\right)(c_L + d) - (t_1 - \tilde{t})K \ \geq \ 0 & \qquad (3.3)
\end{aligned}
$$

for all $\tilde{t} \in [0, t_1)$. The first term refers to the potential gain from setting $t < t_1$ if the athlete is type L. This term may be positive or negative depending on the exact parameter values. The second term refers to the ADA's gain from setting t_1 if the athlete is type H. This term is unambiguously positive.

As was the case when consumers could not observe negative tests, the ADA only benefits from deviating if the athlete is type L and the first term of condition 3.3 is negative. If a type L athlete cheats and either tests negative or is not tested but consumers participate, the ADA gains d by deviating. However, if the athlete tests positive or is not tested and consumers do not participate, the ADA loses c_L. In addition, if the athlete tests positive, the ADA loses B. If type L does not cheat, the ADA's payoff from type L is the same regardless of whether or not it deviates. If the athlete is type H and the ADA deviates, the ADA loses $c_L + d$ if the athlete is not tested and consumers then do not participate.

The left hand side of condition 3.3 is strictly decreasing in $\sigma_{NT}(\tilde{t})$. $\sigma_{NT}(\tilde{t})$ is increasing in F and decreasing in b, so as was the case when consumers could not observe negative tests, an increase in the punishment for an athlete who tests positive actually reduces the range of parameter values for which a no-cheating equilibrium occurs. The left hand side of condition 3.3 is increasing in B for $\tilde{t} > 0$. Therefore, this condition will be satisfied for all $\tilde{t} > 0$ as long as B is sufficiently large. However, when $\tilde{t} = 0$, B does not enter into this condition. When $\tilde{t} = 0$ and as $K \to 0$, the condition instead reduces to:

$$\hat{\gamma} \leq \frac{b/(a_L + b)}{d/(c_L + d)}$$

which is exactly the same condition which was required when consumers could not observe negative tests. Once again, given the set-up, if a no-cheating equilibrium exists, it is also the unique equilibrium.

3.6.2.1 *Proportional Payoffs*

Again, it is possible to conduct further analysis if the simplifying assumption is made that $c_L = \eta a_L$, $d = \eta b$ and $B = \eta F$. As $K \to 0$, the ADA's profit if it sets t in the range $[0, t_1)$ is then:

$$(1-\mu)c_L + \mu \left[t + (1-t)\sigma_{NT}(t) \right] (c_L + d)$$

which as is shown in appendix B.4.1 is increasing in t. The ADA will therefore set $t = t_1 - \epsilon$ if it decides to deviate. If the athlete is type L, the ADA receives the same payoff regardless of whether or not it deviates. The intuition for this is exactly the same as when consumers could not observe negative tests. Once again though, if the athlete is type H and the ADA deviates, its payoff is strictly less than when it does not deviate since if the ADA deviates and the athlete is not tested, consumers may not participate. Therefore, when payoffs are proportional the ADA always prefers to set $t = t_1$.

When payoffs were proportional and consumers could not observe negative tests, it was shown that the ADA would not wish to deviate if $min\{t_1, \hat{t}\} = t_1$, but would wish to deviate if $min\{t_1, \hat{t}\} = \hat{t}$. Therefore it is advantageous to have consumers able to observe negative tests when payoffs are proportional since a no-cheating equilibrium always exists in this case while it may not when negative tests are not observable. When negative tests are not observable, consumers may be willing to participate with certainty if they do not observe a positive test, even when $t < t_1$. The ADA then wishes to deviate. However, when consumers can observe negative tests, the ADA cannot guarantee they will participate unless it sets $t \geq t_1$ and the ADA then has no incentive to deviate. This suggests that making negative, as well as positive, tests observable to consumers would help to eradicate doping in sport.

3.7 EXTERNAL ENFORCEMENT AGENCY

In some cases it may be unreasonable for consumers to be able to observe or understand the probability with which an ADA tests athletes. In this case, it may be necessary for an external enforcement agency to intervene in order to incentivise the ADA to carry out a sufficient number of tests that the athlete is deterred from cheating. A natural candidate for such an external enforcement agency is the World Anti-Doping Agency (WADA). It is assumed that WADA wants to force the ADA to test with a sufficiently high probability that type L will prefer not to cheat.

WADA is able to randomly carry out tests on athletes. The probability with which WADA tests is τ. This probability is announced by WADA at the start of the game so WADA cannot condition its testing probability on whether or not the athlete tests positive. If WADA finds that an athlete has cheated when the ADA found that this athlete had not cheated, the ADA is fined D. If the ADA did not carry out a test on an athlete WADA takes this as meaning that the athlete concerned was found not to have cheated by the ADA. It is assumed that ρ is the same for both the ADA and WADA. If an athlete does not cheat, there is no probability that WADA finds the athlete cheated. If the athlete cheats the probability that the ADA finds that this athlete did not cheat while WADA finds the athlete did cheat is $\rho\tau(1 - \rho t)$.

It is assumed that the athlete is only fined if he tests positive in the test carried out by the ADA. The tests carried out by WADA are purely for the purpose of checking on the ADA and not for checking on the athlete. Therefore, the behaviour of both types of athlete is unchanged and if type L expects consumers to participate, he once again prefers to cheat unless $t \geq t_1$. Type H will never cheat regardless of the t chosen by the ADA and whether or not consumers participate.

It makes no difference whether or not consumers observe the value of τ chosen by WADA since consumers know WADA wishes to deter cheating and therefore can always correctly infer the value of τ chosen by WADA even if they do not directly observe it. However, it is assumed that consumers must make their decision prior to observing the outcome of any test carried out by WADA and also do not observe the testing probability chosen by the ADA. As in the baseline case they observe positive tests but in the absence of a positive test, they cannot tell whether the athlete tested negative or was not tested at all. Once again simultaneous and sequential timing are considered.

3.7.1 *Simultaneous Timing*

The timing of this game is as follows:

1. WADA announces the probability, τ, with which it will test the athlete.
2. The athlete decides whether or not to cheat. Simultaneously the ADA decides the probability with which it will test the athlete.
3. Consumers decide whether or not to participate after observing P and if the athlete tested positive in any test carried out by the ADA.
4. WADA determines whether or not the athlete will be tested, in accordance with the value of τ which it announced at the start of the game. If the athlete tests positive in WADA's test when the ADA failed to detect cheating, the ADA is fined.

Proposition 3.6. *When timing is simultaneous, a no-cheating perfect Bayesian equilibrium does not exist.*

Proof. See appendix B.5. □

As in the baseline case, the actions which would be required on the equilibrium path for a no-cheating equilibrium are:

- ADA: $t = t_1$
- Type H athlete: Don't cheat
- Type L athlete: Don't cheat
- Consumers: Participate

Therefore, once again a no-cheating equilibrium exists provided it is optimal for the ADA to set $t = t_1$. However, the ADA recognises that if it deviates and sets $t = 0$, neither consumers nor the athlete are able to observe this deviation and consequently take the same actions as when the ADA chose $t = t_1$. Since testing is costly, the ADA then prefers to set $t = 0$. When timing is simultaneous, the threat of the external testing authority uncovering doping does nothing to deter the ADA from not testing since the ADA does not expect either type of athlete to cheat if it deviates from $t = t_1$.

3.7.2 *Sequential Timing*

The timing of this game is as follows:

1. WADA announces the probability, τ, with which it will test the athlete.
2. The ADA informs the athlete of the probability with which he will be tested.
3. The athlete decides whether or not to cheat.
4. Consumers decide whether or not to participate after observing P and if the athlete tested positive in any test carried out by the ADA.
5. WADA determines whether or not the athlete will be tested, in accordance with the value of τ which it announced at the start of the game. If the athlete tests positive in WADA's test when the ADA failed to detect cheating, the ADA is fined.

Proposition 3.7. *A no-cheating perfect Bayesian equilibrium exists with sequential timing provided D is sufficiently large.*

Proof. See appendix B.6. □

The equilibrium strategies are as follows. Actions taken on the equilibrium path are highlighted in bold:

- WADA: **Announce $\underline{\tau}(0)$ if $\rho\tau D < c_L + d + B$; announce $\underline{\tau}(t_1 - \epsilon)$ if $\rho\tau D \geq c_L + d + B$.**[4]

4 Either of these sets of parameter values could occur on the equilibrium path.

- ADA: $\boldsymbol{t = t_1}$
- Type H athlete: **Don't cheat**
- Type L athlete: **Don't cheat if $\boldsymbol{t \geq t_1}$**; cheat if $t < t_1$
- Consumers: **Participate**

$\underline{\tau}(0)$ and $\underline{\tau}(t_1 - \epsilon)$ are derived in the appendix. The existence of a perfect Bayesian no-cheating equilibrium again rests on whether or not it is optimal for the ADA to set $t = t_1$. When timing is simultaneous, an external authority does nothing to deter the ADA from deviating and setting $t = 0$ since the ADA does not expect either type of athlete to cheat when it deviates. However, when timing is sequential, the ADA recognises that its choice of t determines the decision of type L about whether or not to take drugs. If the ADA announces $t < t_1$, type L will cheat. There is then a risk of type L testing positive, either in the ADA's own test or in WADA's test. In the baseline case, the ADA could ensure that type L would not test positive even if he cheats by setting $t = 0$. However, with the introduction of additional testing by an external authority, this is no longer possible. Even if the ADA does not test the athlete, he can still test positive in the external's test and the ADA is then fined. Provided D, the fine the ADA must pay if type L tests positive in the external's test, is sufficiently high, the external can always incentivise the ADA to set $t = t_1$ by setting a high value of τ.

3.8 DISCUSSION AND CONCLUSION

The analysis in this chapter has shown that if anti-doping agencies are not benevolent, the current system cannot produce drug-free sport. Anti-doping agencies currently lack the incentives to test athletes with a sufficiently high probability that they will not cheat. Consequently the system must be redesigned if drug-free sport is to become a reality.

The current system corresponds to the baseline case with simultaneous timing. Consumers are able to observe the performance of athletes and positive tests, but cannot observe the frequency with which athletes are tested. In addition, anti-doping agencies currently do not provide athletes with an indication of how often they will be tested. The results presented in this chapter show that unless athletes are informed about the frequency with which they will be tested, a no-cheating equilibrium cannot be achieved, even if the information structure of consumers is changed or there is an external testing authority. The problem with simultaneous timing is that both the anti-doping agency (ADA) and consumers then take the athlete's decision as given and the ADA then always wishes to deviate to not testing at all.

However, informing athletes about testing is not sufficient to achieve a no-doping equilibrium since in the baseline case, even when timing is sequential, a no-doping equilibrium does not exist. When the frequency of testing is not observable to consumers, they cannot condition their decision about whether or not to participate on testing. The ADA then always wishes not to test since it then guarantees the athlete will not test positive. There are consequently two different possible solutions to the problem of doping in sport. The first is to make consumers better informed so that they can condition participation on testing. The second is to introduce an external testing authority so that the ADA can never guarantee a cheating athlete will not test positive, even if the ADA itself does not test the athlete.

The first change to the status quo which was examined was to allow consumers to observe the probability with which the athlete is tested. The World Anti-Doping Agency currently publishes statistics on the number of tests carried out by each country and each sport, but figures are not provided on the number of athletes tested or eligible to be tested. Therefore the simplest possible change would be to force anti-doping agencies to publish statistics on the average frequency with which athletes are tested. If the cost to the ADA from a positive test is sufficiently high and consumers only tolerate a low probability of cheating, a no-cheating equilibrium then exists and this is the unique equilibrium. However, if consumers do not observe negative tests, it is hard for them to verify that the ADA is indeed testing with the frequency it claims. Therefore, the analysis was extended to allow consumers to also observe negative tests. Once again the unique equilibrium involves no-cheating provided the punishment for the ADA after a positive test is sufficiently high and consumers are only prepared to participate if the probability of doping is low. If the athlete's payoffs are proportional to the ADA's, a no-cheating equilibrium always exists when negative tests are observable while this is not the case when negative tests are not observable.

If anti-doping agencies are not willing to publish statistics on the frequency of testing, a no-cheating equilibrium can still be achieved provided there is an external testing authority which can randomly re-test athletes and punish the ADA if it failed to detect doping. These random tests ensure the ADA cannot guarantee the athlete will not test positive, unless it tests with a sufficiently high probability that neither type of athlete wishes to cheat.

There are several possible extensions to the model presented here. Firstly, this model was a one-shot game in which the athlete made a one-off decision about whether or not to cheat. However, it would be desirable to develop a dynamic model in which the athlete could change his decision about whether or not to take drugs and update his beliefs about the probability with which he will be tested. Secondly, there were only two types of athlete in this model. With a continuum of athlete types, the analysis would be more complicated since different types would require different probabilities of testing in order to be willing not to cheat. Finally, in this chapter the punishment for an athlete who tested positive was a lump sum fine. This is the approach taken throughout the existing literature but in chapter 2 it was shown that no-doping may not be optimal when the sanction is a ban for competition. It would therefore be insightful to combine the analysis of chapters 2 and 3 into a single model.

4 TESTING THE TESTERS: DO MORE TESTS DETER ATHLETES FROM DOPING?

4.1 INTRODUCTION

Doping scandals have become increasingly common in sport in recent decades. There exists a growing literature in sports economics which uses theory to analyse this problem. However, far less attention has been devoted to an empirical analysis of the performance of anti-doping agencies. This chapter examines whether testing athletes more frequently deters them from taking drugs. This is an important question since if testing does deter doping, this suggests a relatively simple solution to the problem of doping in sport. If instead testing does not act to deter athletes from taking drugs, we may want to consider redesigning the system.

Ideally we would directly examine the relationship between the number of times an athlete is tested and whether or not the athlete chooses to cheat. However, such an approach is not possible since we do not know whether or not an athlete chose to take drugs. Consequently a less direct approach is required. This chapter uses the relationship between testing and Olympic performance to infer the relationship between testing and doping. This requires a variety of assumptions, the most important of which is that doping improves Olympic performance. This appears to be true of many, although not all, Olympic sports. If doping improves Olympic performance, we would expect that accounting for other factors, countries which test more would perform worse at the Olympics. This is because if testing acts to deter athletes from taking drugs, then countries which carry out more tests will have fewer competitors taking drugs and, if drugs confer a competitive advantage, will subsequently perform worse than would otherwise be expected.

The results suggest that in some sports, such as athletics and wrestling, carrying out more tests does deter athletes from taking drugs. In other sports where doping is believed to be common though, there is no evidence of a negative relationship between testing and doping. This is notably the case in cycling which in the past has been subject to various doping scandals.

This is one of the first empirical papers to analyse the existence of a relationship between doping and testing. The analysis is possible since the World-Anti Doping Agency (WADA) has recently begun to publish detailed statistics concerning the number of tests carried out by National Anti-Doping Organisations (NADOs). To the best of my knowledge, this is the first economics paper to use this new dataset. Two other papers have analysed similar problems although they have taken very different approaches. Hermann and Henneberg (2013) analyse how many tests would be necessary to deter athletes from doping. Using variables such as the window of detection for different drugs and how predictable testing is, they find that depending on the sport between 16 and 50 tests would need to be carried out on athletes each year in order to completely deter doping. Secondly, Mitchell and Stewart (2004) examine whether the announcement that EPO testing would be introduced at the 2000 Olympics led some athletes to choose not to participate. They find no evidence that this was the case.

This chapter also draws on the existing literature which focuses on establishing the factors which best predict how well a country will perform at the Olympics (for example Leeds and Leeds, 2012 and Pierdzioch and Emrich, 2013). These papers examine the importance of a variety of different factors and generally find previous performance, population and GDP per capita to be among the most important variables. Sterken and Kuper (2003) find that communist countries perform better at the Olympics. They hypothesise that this could be the result of state-sponsored doping. However, to date, no paper has included data on testing in an analysis of the determinants of Olympic performance.

4.2 MODEL

If athletes are expected utility maximisers then an athlete will choose to take drugs if his expected utility from doing so exceeds that from not. Let a be an athlete's baseline performance without drug taking which is drawn from an underlying distribution. b is the multiplicative gain which this athlete derives from doping where b is also drawn from an underlying distribution and b is independent of a. a and b are known to the athlete but the anti-doping agency only knows the underlying distributions from which they are drawn. An athlete is tested with probability t and if he tests positive he is banned for proportion n of his career. If an athlete takes drugs he tests positive with probability ρ if he is tested while if he does not take drugs he tests positive with probability θ if he is tested. An athlete is more likely to test positive if he takes drugs so $\theta < \rho$.

If an athlete does not take drugs, with probability $1 - \theta t$ he does not test positive and obtains utility a. With probability θt he tests positive and obtains utility $(1 - n)a$. If an athlete dopes, with probability $1 - \rho t$ he does not test positive and obtains utility ab. With probability ρt he tests positive and obtains utility $(1 - n)$ab. Consequently an athlete dopes if:

$$(1 - \theta nt)a \quad \leq \quad (1 - \rho nt)ab$$

Since $\theta < \rho$, this condition is only ever satisfied if $b \geq 1$. When $b \geq 1$, the athlete will dope if:

$$t \quad \leq \quad \frac{b-1}{n(pb-\theta)} \equiv \beta(b)$$

If t increases while n, the punishment for testing positive, is held constant, $Pr\left[t \leq \beta(b)\right]$ decreases, as long as ρ, the probability of a true positive is strictly positive. Consequently, as the probability with which an athlete is tested increases, provided the punishment for testing positive remains unchanged, we would expect the proportion of athletes taking drugs to decline. It is important to note that we would not necessarily expect to see a negative relationship between testing and doping if countries were able to compensate for a low testing probability by setting a harsher punishment. However, in reality countries cannot do this since WADA's Code (WADA, 2015*b*) specifies set punishments which are standardised across sports and all countries which competed at the 2012 Olympics are signatories to the Code.

Even with standardised punishments though, an increase in the testing probability only works to deter an athlete from doping if it increases the probability with which he expects to test positive if he takes drugs, ρt. Consequently, if ρ and θ are both equal to zero, an increase in t does not deter drug taking because the athlete expects that even if he takes drugs and is tested, he will never test positive.

4.2.1 *Ideal Approach*

Ideally we would investigate the relationship between testing probability and an athlete's choice of doping by carrying out a randomised controlled trial in which different athletes were tested with different probabilities and we could observe whether or not each athlete chose to take drugs. However, even assuming that permission for such a trial could be acquired from an anti-doping agency, it would not be possible to accurately observe whether or not an athlete had taken drugs. Consequently it is necessary to use a more indirect approach to analyse this issue.

4.2.2 *Alternative Approach*

The approach this chapter takes is to use the relationship between testing and Olympic performance to infer a relationship between testing and doping. The relationship I am really interested in is how the proportion of athletes who choose to dope in country i and sport s, D, changes as samples per athlete carried out by country i in sport s, SPA, changes.[1] The

1 For notational simplicity, the subscripts i and s are omitted from all variables in this section.

proportion of athletes who choose to dope is a function of the frequency of testing, SPA, and other variables, denoted by the vector x. Therefore, $D(SPA, x)$. However, we do not observe $D(SPA, x)$, and therefore cannot directly calculate $\partial D/\partial SPA$ since we do not know the proportion of athletes choosing to dope.

The proportion of medals won by country i in sport s is denoted by M. This is a function of doping and other variables so $M(D, x)$. We also do not observe $M(D, x)$ since doping is not known. However, it is assumed that once other variables are taken into account, testing only affects the proportion of medals via doping. This is an exclusion restriction. This assumption is reasonable since there is no reason to believe that testing would directly affect the number of medals won once factors such as funding for sport have been taken into account. We can therefore substitute $D(SPA, x)$ into the medal function to obtain:

$$M\Big(D(SPA, x), x\Big)$$

Using μ to denote medals as a function of SPA and x, we then have:

$$\mu(SPA, x) \quad = \quad M\Big(D(SPA, x), x\Big)$$

Then, by the chain rule:

$$\underbrace{\frac{\partial \mu(SPA, x)}{\partial SPA}}_{1} = \underbrace{\frac{\partial M\Big(D(SPA, x), x\Big)}{\partial D}}_{2} \underbrace{\frac{\partial D(SPA, x)}{\partial SPA}}_{3}$$

Term 3, the change in the proportion of athletes doping when testing changes, is the term we are really interested in but data is not available to directly calculate term 3. Instead, data is available to calculate term 1, the change in Olympic performance when testing changes. With the exclusion restriction, any change in testing only affects the medal proportion via doping once other factors have been taken into account. From the chain rule it can be seen that it is then possible to infer the sign, although not the magnitude, of term 3 from term 1 provided term 2 can be signed.

Term 2 is the change in Olympic performance when the proportion of athletes doping changes. There is a sizeable body of scientific evidence suggesting that doping improves performance in a wide range of sports such as cycling and weightlifting. In a study on recreational athletes, taking human growth hormone resulted in a 4% increase in sprinting capacity while in another, blood doping increased stamina by 34% (Thompson, 2012). Consequently, we would expect that for at least some Olympic sports, there would be a positive relationship between the proportion of athletes from a country who dope and Olympic performance of the country in question. This implies that, especially for sports with a sizeable strength, stamina or speed component, $\partial M / \partial D$ is likely to be positive. However, for other sports, and in particular skill based sports, there is little evidence that drugs are able to improve performance.

If term 2 is positive, this means that the sign of term 1 must be the same as the sign of term 3. Consequently, if we find a negative relationship between Olympic performance and testing, this implies the relationship between testing and doping must also be negative.

4.3 EMPIRICAL FRAMEWORK

4.3.1 *Data*

The dependent variable is the proportion of available medals won by country i in sport s at the 2012 Olympic Games. The source for this data was the official website of the Olympic Games.[2] All medals were weighted equally so regardless of whether a country won gold or bronze, this was recorded as one medal. Using medal counts or medal shares is the standard approach used in the literature on Olympic performance by authors such as Bernard and Busse (2004), Johnson and Ali (2004), Leeds and Leeds (2012) and Pierdzioch and Emrich (2013). There are several reasons for using medal shares as my measure of Olympic performance. Firstly, since at the elite level, even a small performance improvement can make the difference between winning a medal or not, we might

2 http://www.olympic.org/

expect medal shares to be particularly responsive to doping. Secondly, if a weighting system were to be employed, it is not clear how a gold medal should be weighted relative to silver or bronze. Finally, if gold, silver and bronze medals were analysed separately, the number of observations which do not have a positive entry for the dependent variable would increase even further. It was necessary to normalise the medal counts by the number of available medals in each sport since different sports award different numbers of medals.

In total there are 1673 observations at the level of country and sport. There are 204 countries in the data set and 26 categories of sport. Sports in which fewer than 30 countries had at least one athlete competing are categorized as 'Other' for the purposes of the analysis.[3]

Category of Sport	Number of Countries Competing
Aquatics[4]	168
Archery	55
Athletics	200
Badminton	51
Basketball	17
Boxing	79
Canoe	56
Cycling	74
Equestrian	40
Fencing	44
Football	24
Gymnastics[5]	57

3 Basketball, Football, Handball, Hockey, Modern Pentathlon.
4 Diving, Swimming, Synchronized Swimming, Water Polo
5 Gymnastics Artistic, Gymnastics Rhythmic, Trampoline

Handball	17
Hockey	15
Judo	134
Modern Pentathlon	26
Rowing	58
Sailing	62
Shooting	108
Table Tennis	57
Taekwondo	63
Tennis	44
Triathlon	39
Volleyball	30
Weightlifting	84
Wrestling	71

Table 4.1: Categories and number of countries competing

The source of data for testing statistics is the World Anti-Doping Agency. All data on testing statistics is presented in the 2012 Anti-Doping Testing Figures Report which is publicly available on WADA's website (WADA, 2013*a*).[6] This data was published for the first time in relation to testing carried out in 2012. The report provides a detailed breakdown of the number of tests carried out by each country in each sport. For example, data is available on the number of tests carried out by the United Kingdom on cyclists. Since several organisations within the same country may carry out drug tests, the figures for each organisation within a country were aggregated in order to produce a single figure for each country. In some cases, WADA only provides statistics for a category of sport such as aquatics, rather than for the individual constituent sports such as swimming and

6 https://www.wada-ama.org/

diving. In this case, medal proportions were calculated by aggregating medals across all sports in the category.

Some countries, and in particular many third-world countries, lack the resources to conduct their own testing programme. Such countries are often members of a Regional Anti-Doping Organisation (RADO). RADOs pool resources in order to carry out tests on athletes belonging to member countries.[7] Information is only available on the total number of tests carried out by each RADO in each sport; we do not have data on the number of tests carried out by the RADO on athletes of each member country. Therefore RADO tests were allocated to member countries in proportion to the size of the team which each country sent to the Olympics. Since in total only 403 tests were carried out by all RADOs across all sports, this step is unlikely to significantly impact on the results.

The difficulty with this data set is that data is only provided at the level of country and sport. There is not data on how many times an individual athlete was tested or even on how many athletes in total were tested. Consequently, there is no information about the number of cyclists the United Kingdom tested, only the total number of tests carried out on cyclists. However, I am interested in how doping varies with tests per athlete, not total tests. The solution I adopt is to divide testing statistics by the team size country i sent to the 2012 Olympics for sport s. If for example the UK carried out ten tests on cyclists and sent two cyclists to the Olympics, we would record that the UK carried out five tests per cyclist. Data on team size was acquired from the website of The Guardian,[8] a British newspaper.

We might also expect that Olympic performance would display a strong element of path dependence with countries which previously performed well at certain sports also performing well in 2012. Therefore the proportion of medals won by country i in sport s in 2008 is included as an independent variable. This captures the ability of a country in a specific sport.

---

7 For example, the participating countries in the South Asia RADO are Bangladesh, Bhutan, Maldives, Nepal and Sri Lanka.

8 http://www.theguardian.com/sport/datablog/2012/jul/30/olympics-2012-alternative-medal-table#data

In order to account for the possibility that a variety of omitted variables may affect both testing and Olympic performance, country fixed effects are included in the regression. Since variables such as GDP p.c., population and corruption do not vary by sport, the effect of these variables is also captured through the fixed effects.

Summary statistics for the variables pooled across all observations are given in table 4.2.

Variable	Mean	Std. Dev.	Min	Max
2012 medal proportion	0.0155409	0.0450829	0	0.5333334
Samples per athlete	7.549979	17.61526	0	414
2008 medal proportion	0.0150872	0.0472328	0	0.6666667

Table 4.2: Summary Statistics

4.3.2 *Potential Issues*

With an ideal dataset, it would be possible to infer the sign of the relationship between testing and doping from the sign of the relationship between testing and Olympic performance. However, due to a lack of data, there are several issues which may hamper any such attempt. These issues can be divided into three categories: issues of aggregation, issues of omitted variables and issues of timing.

4.3.2.1 *Issues of Aggregation*

The first potential issue can be seen from table 4.2 which shows that the maximum number of tests per athlete in the pooled dataset is 414. This figure is clearly implausibly high and occurs due to my method of handling the problem that the finest division of testing statistics is at the country-sport level. Since I only have data on the total number of tests carried out by a country in a sport, I created a value for samples per athlete by dividing by Olympic team size. The difficulty with this approach is that not all athletes who are tested may be sent to the Olympics. In a small number of cases, this results in a value for samples per athlete which is infeasibly high. Therefore some of the regressions were run restricting observations to those where samples per athlete was less than once per week.

The second issue of aggregated test statistics is that even if the Olympic team size accurately conveyed the number of athletes being tested, average figures may mask effects if testing is targeted. For example, we can imagine a simple case where there exist two heterogeneous groups of athletes. Group A may dope if the probability with which they are tested is sufficiently low. Group B has high moral standards and will never take drugs regardless of the testing probability. An anti-doping agency moves from a regime in which all athletes are tested with equal probability to one in which athletes in group A are tested more often than athletes in group B but the total number of tests is unchanged. The result of this would be a decline in average performance, but no change in total tests. Consequently, even though testing is effective (it deters group A from doping), this is not captured in the relationship between the average number of tests per athlete and performance. This model is only able to assess whether carrying out a greater number of tests deters cheating, it is incapable of analysing whether an improvement in the quality of testing, for example through targeted testing, with no change in the total number of tests carried out, acts as a deterrent. This is because no data is available on the quality of testing by each national anti-doping agency.

Consequently, in this model, quality of testing is an omitted variable and may result in omitted variable bias. This issue will arise if countries have different qualities of testing and the quality of a country's testing is systematically related to the number of tests which are carried out. The direction of the bias depends on whether quality and number of tests are positively or negatively correlated. It seems most likely that quality and number of tests would be positively correlated since countries with greater resources are likely to be able to both carry out more tests and devote more money to better targeting tests. In this case, the relationship between number of tests and performance would be biased in a negative direction. Alternatively, there could be a negative relationship between the number of tests performed by a country and the quality of testing. This would occur if countries compensated for a lack of quality testing by carrying out a greater number of tests. In this case, any negative relationship between number of tests and performance could be masked by a correspondingly lower quality of testing positively impacting on performance. Most countries have only one organisation which is responsible for carrying out tests for all sports. Consequently, we would not expect the quality of testing carried out by a country to differ systematically across sports. Therefore including country fixed effects in the regression should account for cross country differences in quality of testing.

4.3.2.2 *Issues of Omitted Variables*

We can only interpret the coefficient on samples per athlete if there is not omitted variable bias. This will occur if there is a variable which is correlated with both SPA and Olympic performance which is not accounted for. An obvious candidate for such a variable is funding for sport. Countries with more funding for sport are likely to perform better at the Olympics since they are able to spend more money on coaches and facilities. This funding is also likely to be used for testing so these countries also carry out more tests. If this is true and this effect is not taken into account, the coefficient on samples per athlete would be positively biased, since it would also be capturing the effect of increased funding. Since data is not available on funding for sport, country fixed effects are included in the regression.

The second potential issue with regard to omitted variables occurs because the data on testing only captures the average number of times an athlete was tested by their national anti-doping agency. However, athletes can also be tested by other organisations such as the testing authority for their sport. Although data is published on how many tests were carried out by other organisations in a sport, there is no information regarding how these tests are distributed across countries. In an ideal data set, we would observe the number of times an individual athlete was tested by any anti-doping agency, SPA_{Tis}. Using SPA_N to refer to samples per athlete conducted by a national anti-doping agency and SPA_O to refer to samples per athlete carried out by other organisations, and for notational simplicity, ignoring other variables, we would want to estimate:

$$\begin{aligned} M_{is} &= \alpha + \beta_1 SPA_{Tis} + u_{is} \\ &= \alpha + \beta_1 SPA_{Nis} + \beta_1 SPA_{Ois} + u_{is} \end{aligned}$$

Instead what we are estimating is:

$$M_{is} = \alpha + \gamma_1 SPA_{Nis} + v_{is}$$

Therefore:

$$\text{plim}\hat{\gamma}_1 = \beta_1(1 + \text{plim}\hat{\delta})$$

where $\hat{\delta}$ is the coefficient from a regression of SPA_{Ois} on SPA_{Nis}. It is possible that if a country carries out more tests on its athletes, other testing agencies choose to carry out fewer tests. If this is the case SPA_{Nis} and SPA_{Ois} would be negatively correlated, implying $\hat{\delta} < 0$. It is not clear to what extent coordination of testing between authorities does occur but regardless, we would not expect that when a country carries out one more test on an athlete, other organisations would reduce their number of tests by more than one. If it did in fact lead to a reduction of more than one, this would imply that the net effect of a country carrying out one more test would be a reduction in the number of total tests which seems implausible. Therefore, $-1 < \hat{\delta} < 0$ and the sign of $\hat{\gamma}_1$ will be correct. However, we have attenuation bias since $\hat{\gamma}_1$ is biased towards zero. The regression results will accurately convey the impact of an increase in national testing on Olympic performance. However, this may provide an under-estimate of the impact of an overall increase in testing if there is negative correlation between the tests carried out by different authorities.

4.3.2.3 *Issues of Timing*

Finally, issues of timing occur as a result of data availability. The use of testing statistics for the whole of 2012 is potentially problematic, but unfortunately WADA does not publish testing statistics by month, nor were they published prior to 2012. The difficulty with using testing statistics from 2012 to analyse Olympic performance in 2012 is that while the Olympics took place in August, the testing year ran until December. Ideally I would have data on the number of samples per athlete in the year prior to the Olympics. Instead some of the samples recorded in the testing statistics were carried out after the Olympics. This would be of particular concern if there was feedback from Olympic performance to testing. Fortunately, it seems unlikely that such feedback would occur. Firstly, it is unlikely that sports' budgets would be reallocated so quickly after the Olympics. Secondly, in most countries testing is carried out by a national testing organisation, not a national sports specific organisation. It seems most likely that the budget for such an organisation would be independent of Olympic performance.

Even without feedback, the use of data from the entirety of 2012 rather than the year prior to the Olympics introduces measurement error into the testing statistics. Using the classical errors in variables model, coefficients will be biased towards zero. However, the variance of the measurement error is likely to be small and as a robustness check, results were also derived using data from the 2013 Athletics World Championships.

It is also possible to use the regression results to determine the extent to which problems have arisen due to normalising by Olympic team size and funding being an omitted variable. Unfortunately though, without a more detailed breakdown of the testing statistics it is not possible to assess the extent to which the other potential problems have occurred.

4.3.3 *Regression Specification*

The dependent variable in the regression is a proportion. The ideal would therefore be to carry out a non-linear regression such as GLM with family set to binomial and link to logit. However, the problem with such non-linear models is that the coefficients are generally biased when fixed effects are included in the regression (Greene, 2004). There is considerable controversy surrounding whether or not it is appropriate to use the linear probability model when the dependent variable is limited. In the context of this chapter, the advantage of using OLS is that including fixed effects then does not bias the coefficients. However, OLS can produce predictions which are less than zero or greater than one. Given there is no clear consensus on the best approach in this situation, regressions were carried out using both OLS and GLM.

The regression specification is:

$$12M_{is} = \alpha + \sum_{s=1}^{s=22} d_s[\beta_{1s}SPA_{is} + \beta_{2s}8M_{is}] + \gamma_i + \epsilon_{is}$$

where $12M_{is}$ is the proportion of medals won by country i in sport s in 2012 and $8M_{is}$ is defined analogously for 2008. SPA_{is} is the number of samples per athlete carried out by country i in sport s in 2012. γ_i are the country fixed effects. d_s is a dummy variable equal to one for sport s and zero otherwise. The standard errors are clustered by country.

4.4 RESULTS

	I	II	III	IV
Method	OLS	GLM	OLS	GLM
Observations	All	All	SPA <53	SPA <53
Aquatics	-0.000162*	-0.0500*	-0.000533**	-0.0445
	(8.99e-05)	(0.0296)	(0.000263)	(0.0308)
Archery	-0.000898	-0.0806	-0.000910	-0.0741
	(0.000553)	(0.0567)	(0.000594)	(0.0572)
Athletics	-0.000204**	-0.0574**	-0.000312*	-0.0536**
	(9.40e-05)	(0.0263)	(0.000187)	(0.0271)
Badminton	-0.000721*	-0.0625	-0.000691*	-0.0557
	(0.000384)	(0.0387)	(0.000373)	(0.0371)
Boxing	-0.000142**	-0.0179**	-0.000125	-0.0200

	(7.12e-05)	(0.00879)	(0.000243)	(0.0141)
Canoe	-0.000204	-0.0361*	-0.000223	-0.0341
	(0.000287)	(0.0202)	(0.000296)	(0.0208)
Cycling	-7.79e-05	-0.00575	-8.32e-05	-0.000127
	(7.35e-05)	(0.00608)	(0.000216)	(0.00878)
Equestrian	6.97e-05	-0.0266	8.92e-05	-0.0216
	(0.000904)	(0.0690)	(0.000924)	(0.0712)
Fencing	5.28e-05	-0.0111	3.89e-05	-0.00828
	(0.000416)	(0.0310)	(0.000401)	(0.0297)
Gymnastics	-0.000168	-0.0309	-0.000170	-0.0256
	(0.000186)	(0.0297)	(0.000164)	(0.0257)
Judo	-0.000226**	-0.0383**	-0.000157	-0.0593***
	(9.91e-05)	(0.0169)	(0.000208)	(0.0225)
Other	0.000173	-0.00497	0.000783	0.00894
	(0.000516)	(0.00715)	(0.00103)	(0.0228)
Rowing	-0.000358***	-0.0596***	-0.000846***	-0.0551**
	(0.000108)	(0.0228)	(0.000217)	(0.0234)
Sailing	-0.00147	-0.152	-0.00156	-0.136
	(0.00108)	(0.0990)	(0.00115)	(0.103)
Shooting	-0.000427**	-0.102***	-0.000442**	-0.0972**
	(0.000178)	(0.0393)	(0.000198)	(0.0396)
Table Tennis	-0.000434**	-0.0516	-0.000441*	-0.0467
	(0.000211)	(0.0382)	(0.000242)	(0.0392)
Taekwondo	0.000378	0.00522	2.46e-05	-0.00381
	(0.000231)	(0.00506)	(0.000311)	(0.0176)
Tennis	-0.000232	-0.0143	-0.000272	-0.0108

	(0.000475)	(0.0291)	(0.000488)	(0.0289)
Triathlon	0.000242	0.00169	0.000570	0.0227
	(0.000433)	(0.0119)	(0.000647)	(0.0247)
Volleyball	-0.000243	-0.0452	-0.000341	-0.0350
	(0.000357)	(0.0481)	(0.000766)	(0.0531)
Weightlifting	-7.80e-05**	-0.0328***	-0.000314	-0.0367*
	(3.14e-05)	(0.0107)	(0.000196)	(0.0193)
Wrestling	-0.000379***	-0.0462***	-0.000441***	-0.0577**
	(0.000146)	(0.0176)	(0.000165)	(0.0249)
R^2	0.5268		0.5281	
AIC		485.55		480.0358
Observations	1673	1673	1634	1634

Table 4.3: Coefficients on samples per athlete.
Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.
The regression also included the proportion of medals won in 2008 and country fixed effects.

Results are presented in table 4.3. The t-statistics are summarised in figure 4.1 where the red line represents the critical value for a two-sided test at the 5% significance level. Athletics, judo, rowing, shooting and wrestling all have a coefficient on samples per athlete which is negative and significant at the 5% level in at least three of the regression specifications. Boxing and weightlifting have significant coefficients both when OLS and GLM are used, but only when all observations are included. Aquatics and table tennis both have a significant coefficient in only one of the regressions. These results therefore suggest that while in some sports there exists a negative relationship between testing and doping, in other sports no such relationship exists.

Results are presented in table 4.3 from regressions using both OLS and GLM techniques. While the magnitudes of the OLS and GLM coefficients cannot be compared, their

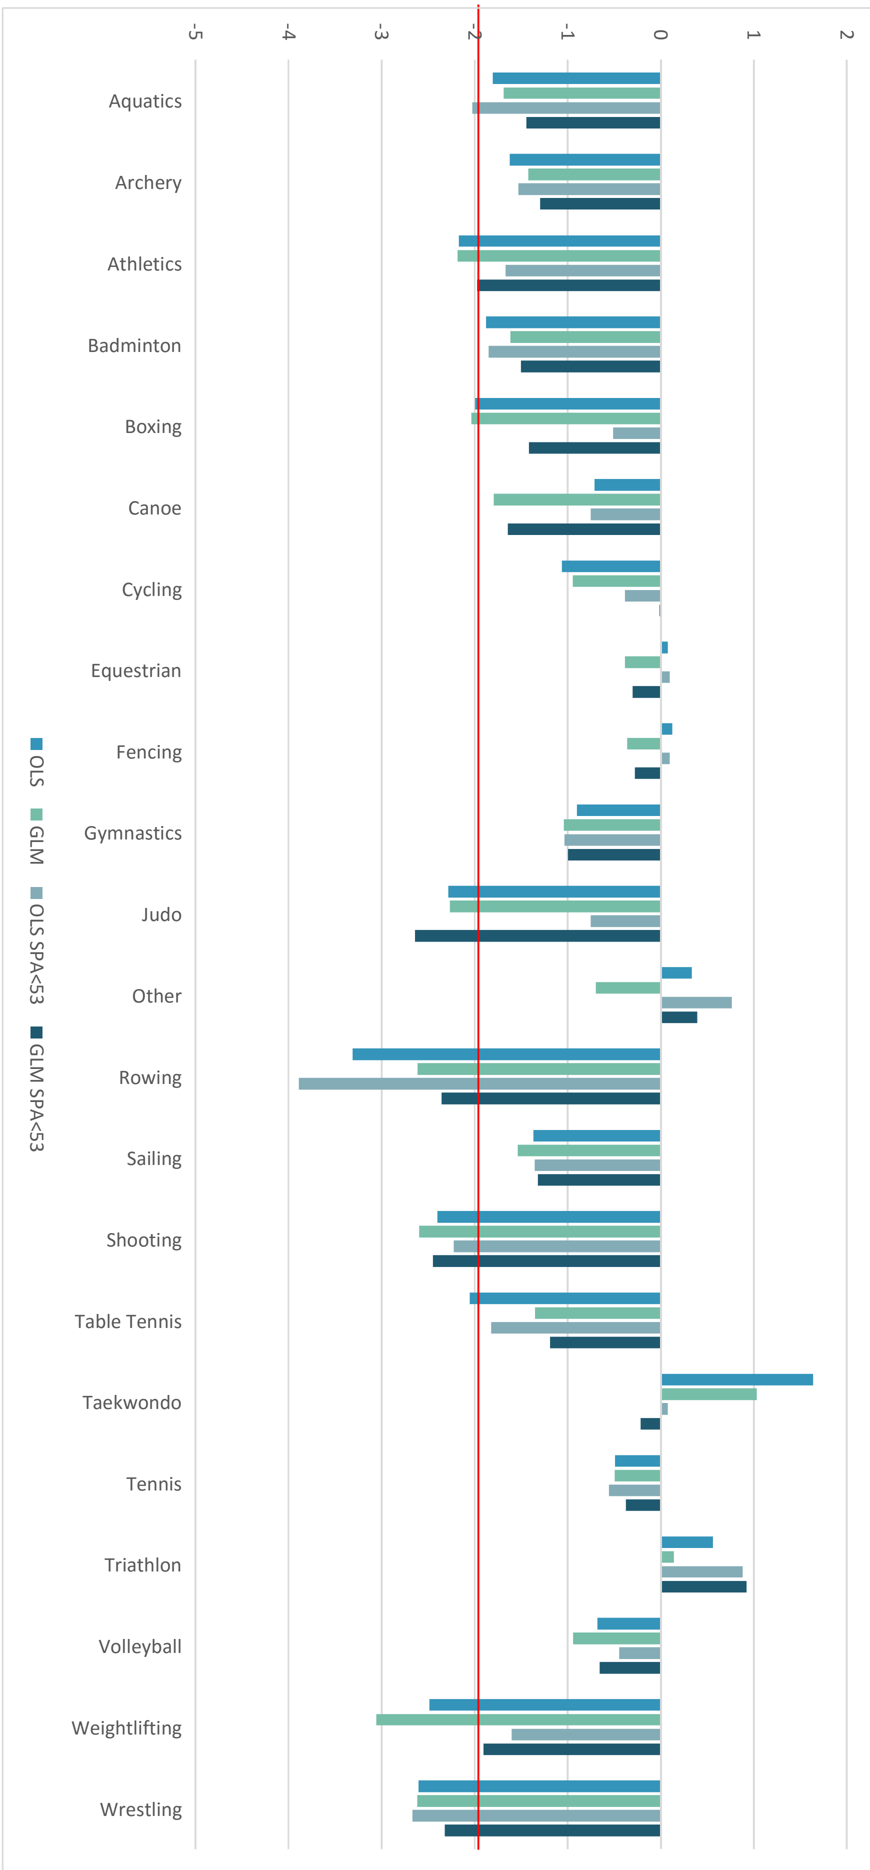

Figure 4.1: t-statistics for the coefficients on samples per athlete

statistical significance can be. This suggests the OLS and GLM techniques produce similar results since for all except four sports, both techniques show either no significance at the 5% level or significance at least at the 5% level. In addition, results are presented both when all observations are included in the regression and when observations are restricted to those where samples per athlete is less than once per week. Restricting observations means that 39 observations are dropped.

The coefficients on proportion of medals won in 2008 are given in table C.1 in appendix C.1. For the vast majority of sports and specifications, the coefficient is positive and highly significant. This confirms the finding of other studies, such as Bernard and Busse (2004), which have shown that Olympic performance is highly path dependent.

4.4.1 *Robustness*

Using these results, it is possible to assess, at least to some extent, the likelihood that the problems discussed above have arisen with this data set. Since many countries have only one national testing agency, as opposed to separate national testing agencies for each sport, we would expect any biases which arise as a result of the available data to affect the coefficients on all sports in a similar manner. The first issue which was discussed was that since the testing statistics are only available at the country-sport level, data on samples per athlete was generated by dividing these statistics by Olympic team size. In a minority of observations though, this approach results in the number of samples per athlete being implausibly high. Consequently, regressions were carried out both on the full set of observations and restricting observations to those where samples per athlete is less than once per week. Table 4.3 suggests that the results are to some extent dependent on observations with high values of samples per athlete. Comparing between the regressions where observations are restricted and those where they are not, in 16 sports, restricting observations does not alter whether or not the coefficient on samples per athlete is significant at the 5% level for both OLS and GLM. Five sports have a coefficient which is significant at the 5% level with all observations but do not when observations are restricted when either

OLS or both OLS and GLM techniques are used. Aquatics is the only sport which has a significant coefficient when observations are restricted, but otherwise does not.

This suggests that for some sports, the analysis is sensitive to the inclusion of high values of samples per athlete, especially when OLS is used. Considering the sports which seem to be sensitive to this suggests that issues of aggregation are likely to be worst in sports where many more athletes compete than there are available places at the Olympics. In addition, sports such as boxing have a strong divide between amateur and professional athletes, with only amateur athletes being eligible to compete at the 2012 Olympics.

Another issue which was discussed was that funding could simultaneously affect testing and Olympic performance. This would result in the coefficients on samples per athlete being biased in a positive direction. It is therefore reassuring that no sport has a positive and significant coefficient. It is likely that in some sports, especially skill based sports, drug taking offers little benefit and therefore does not occur. Such skill based sports may also be precisely those sports where increased spending on coaching and facilities is most beneficial. Consequently, if the analysis was picking up residual effects of funding, we would expect at least some sports, especially sports such as sailing, to have positive and significant coefficients. It is therefore reassuring that none of the regressions have a positive and significant coefficient on samples per athlete for any sport.

A further potential issue is that while the Olympics were in August, the testing year ran until December. In order to check the results are not being biased as a result of feedback from Olympic performance to testing, the above regressions were repeated, substituting the 2012 medal proportions for athletics with data from the 2013 Athletics World Championships. The coefficients on samples per athlete for athletics were comparable in both magnitude and significance suggesting that feedback is not a cause for concern

As an additional robustness check, a panel dataset was compiled for athletics. This consisted of the proportion of medals won at the 2012 Olympics and the 2013 Athletics World Championships as well as samples per athlete by country for both years. Samples per athlete for 2013 were computed using the WADA 2013 testing statistics for total number of tests carried out by each country in athletics, and then dividing by the team size each country sent to the 2013 World Championships. Coefficients were computed using the regression specification[9]:

$$M_{it} = \alpha + \beta_1 SPA_{it} + \gamma_i + \epsilon_{is}$$

where t refers to the time period.

The same four techniques as above were used to compute the coefficient on samples per athlete. The results are given in table 4.4.

	I	II	III	IV
Method	OLS	GLM	OLS	GLM
Observations	All	All	SPA <53	SPA <53
SPA	-0.000124*	-0.0400**	-0.000144**	-0.0400**
	(6.28e-05)	(0.0186)	(6.88e-05)	(0.0186)
R^2	0.0059		0.0186	
AIC		336.69		332.69
Observations	406	406	402	402

Table 4.4: Coefficients on samples per athlete from athletics panel dataset. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. The regression also included country fixed effects.

In the main dataset the coefficient on samples per athlete for athletics was negative and significant in the majority of the specifications. This is also the case in the panel dataset

9 Data is available for performance at the 2011 Athletics World Championships but testing statistics are not available for this year. Since there are three periods of data for the dependent variable, it is also possible to to compute coefficients for a regression specification which includes both fixed effects and lagged performance using the Arellano-Bond estimator. When this is done, the coefficient on samples per athlete is negative but not significant. However, the coefficient on lagged performance is also not significant while in the main regression it was highly significant. This suggests it is not appropriate to include both fixed effects and lagged performance in the regression. In the main dataset which consists of only Olympic data, we can only compute general fixed effects rather than sport specific fixed effects. Consequently it is necessary to also include previous Olympic performance in order to account for the ability of a given country in the specific sport in question. In contrast, in the athletics panel dataset, the fixed effects computed are country fixed effects specifically for athletics. Therefore for the athletics data, fixed effects already capture a country's ability specifically in athletics.

suggesting that the results from the main dataset are indeed robust. We would not necessarily expect the magnitude of the coefficients to be comparable across these two datasets since for the panel dataset fixed effects capture country fixed effects specifically for athletics, while when the main dataset was used, the fixed effects were not sport specific. The values of R^2 are so low for the OLS regressions because this measure neglects the estimated fixed effects.

4.5 DISCUSSION

Overall, the results suggest that in some sports the anti-doping picture is potentially less bleak than it is often portrayed as being. This is particularly true of sports such as athletics, judo, rowing, shooting and wrestling. In addition, boxing and weightlifting perform well in some of the specifications. These are all sports where we would expect doping to potentially improve performance. In endurance athletics events and rowing, methods of doping which improve an athlete's red blood cell count are known to be effective (Thompson, 2012). Meanwhile in disciplines requiring strength or speed, steroids have been shown to enhance performance. At first, the presence of shooting in the list of sports where testing appears to deter doping may be puzzling. However, shooters can benefit from using beta-blockers to reduce tremors allowing them to have a steadier aim (Oransky, 2008).

But in other sports where doping is believed to be prevalent it does not appear to be the case that increasing testing reduces doping. It is notable that cycling does not have a negative and significant coefficient on samples per athlete in any of the regression specifications, despite the fact that blood doping is believed to be widespread in the sport.

These results imply that in sports such as athletics, carrying out more tests may deter athletes from doping. It is important to note that the results do not imply that doping is not an issue in the sports with a negative and significant coefficient. Nevertheless, they do suggest that a relatively simple, though expensive, solution to the doping problem is to test athletes more frequently. In other sports though, such as cycling, it seems that the problem of doping may have deeper roots.

In order to better understand this, we need to consider why an athlete may not change his decision about whether or not to take drugs when he is tested more frequently. The key reason for this is if the athlete does not believe he will test positive regardless of how often he is tested. In turn there are two main explanations for why an athlete may not believe he will ever test positive. The first is if the tests being used by anti-doping agencies are not capable of detecting the substances being used by athletes. For example, a test for THG, a designer steroid, was only developed when the US anti-doping agency was anonymously sent a syringe containing the substance. A second possibility is that anti-doping agencies are colluding with athletes. This would be the case if anti-doping agencies were either warning athletes about upcoming tests or falsifying test results. There is evidence that the Russian Anti-Doping Agency has engaged in both of these practices (Pound et al., 2015). The possibility of a non-benevolent testing agency and potential solutions to this problem were discussed in chapter 3.

However, it is important to treat the results presented in this chapter with caution. Several assumptions were required to conduct the analysis mainly as a result of limited data availability. While it appears that many of the assumptions are indeed satisfied, it is not possible to check the validity of all of the assumptions.

4.6 CONCLUSION

The results presented in this chapter suggest that while for some sports increased testing deters athletes from doping, in other sports there is no evidence that this is the case. Sports such as athletics and wrestling appear to have testing regimes which are potentially capable of deterring doping while cycling does not. This suggests that in sports such as athletics, carrying out more tests would be an effective method for solving the problem of doping. In sports such as cycling though, more analysis is required to determine why athletes may not believe they are at risk of testing positive despite being tested.

As the World Anti-Doping Agency publishes testing statistics for more years, a panel data approach to analysing the relationship between testing and doping will become possible. In particular, after the 2016 testing statistics are published, it will be possible to compile a panel dataset consisting of two summer Olympic Games. In the meantime, WADA continues to make new information available and for the first time, the 2013 testing statistics contain a breakdown of testing in specific disciplines, such as marathon running, within a sport. The availability of more detailed data and a greater quantity of it should enable more in-depth analysis to be carried out in the future.

5 CONCLUSION

This thesis has examined optimal policy responses to the problem of doping in sport. This is a problem which appears to be both escalating and attracting increasing public attention. There is now evidence that not only are athletes doping, but individuals and organisations at the very highest level of sport have been helping to ensure that such cheating goes undetected. The report recently published by an independent commission on behalf of WADA (Pound et al., 2015) has shown that the Russian Anti-Doping Agency was involved in covering up doping by Russian athletes and the International Association of Athletics Federations (IAAF) is being investigated for corruption in conjunction with this. This thesis has aimed to offer practical suggestions as to how such crises can be averted in the future.

Firstly, I analysed the optimal ban for a benevolent organisation to impose on an athlete who tests positive. In the existing literature the punishment is modelled as a lump-sum fine. The novel contribution of this chapter is that the punishment is modelled as a ban rather than a lump-sum fine which reflects the method of punishment used in sport. It has generally been assumed that maximal deterrence from doping is optimal for society, but I have shown that when the punishment is a ban from competition, this may not be socially optimal. This is because when the punishment is a ban, if an athlete tests positive society does not benefit from the athlete's performance for the duration of his ban. This chapter then analysed how the optimal punishment should vary with the attributes of different sports. WADA's current policy is to have punishments standardised across sports. However, my analysis suggested that WADA should consider relaxing this policy since the optimal punishment is found to depend on factors such as the prevalence of doping in the sport and the testing regime.

Secondly, I examined the possibility that the testing agencies responsible for catching cheats may not wish to do so. This analysis is particularly relevant in light of the recent

Russian doping scandal. I found that if consumers cannot observe the frequency with which an athlete is tested, a no-cheating equilibrium does not exist if the ADA is non-benevolent. Two main mechanisms were found which can incentivise such an ADA to test with a sufficiently high probability that athletes are deterred from doping. The first was if consumers can observe the testing frequency and are only prepared to participate if they believe the probability that the athlete cheated is sufficiently low. The second was if there is an external testing authority, such as WADA, which can randomly retest athletes and punish the ADA if it failed to detect doping.

Finally, my empirical chapter investigated whether increased testing does in fact deter athletes from doping. This question cannot be examined directly since there is not reliable data on doping. Instead I used the relationship between testing and Olympic performance to infer the relationship between testing and doping. The analysis suggested that in some sports, such as athletics and wrestling, increased testing does deter doping. This suggests that a simple solution to the problem of doping in these sports is to test more frequently. However, in other sports such as cycling, there is no evidence that increased testing deters doping. Further investigation is required into the causes of this. It may be that the tests are not capable of detecting the substances being used by athletes. Alternatively, as was analysed in chapter 3, it may be the case that testing agencies lack incentives to catch athletes.

The economic analysis of doping in sport is a relatively new field but economics is particularly well suited to finding solutions to the problem of doping. Averting future doping scandals requires not only better drug tests but also redesigning the anti-doping system so that incentives are better aligned. This thesis has suggested some ways in which this could be achieved, but more analysis is required into the reasons why testing may fail to deter athletes from doping. In addition, as more data is collected on athletes, it will become possible to build an econometric model which would predict whether an athlete is doping based on a variety of factors such as performance improvements and scientific tests. Such a model could then inform testing agencies of the best athletes to target for testing, enabling these agencies to use their scarce resources more efficiently.

BIBLIOGRAPHY

Avenhaus, R., Stengel, B. V. and Zamir, S. (2002), Inspection Games, *in* R. Aumann and S. Hart, eds, 'Handbook of Game Theory', Vol. 3, chapter 51.

BBC (2013), 'Lance Armstrong and Oprah Winfrey: cyclist sorry for doping', http://www.bbc.co.uk/sport/0/cycling/21066354.

Becker, G. (1968), 'Crime and Punishment: An Economic Approach', *Journal of Political Economy* **76**(2), 169–217.

Berentsen, A. (2002), 'The economics of doping', *European Journal of Political Economy* **18**(1), 109–127.

Berentsen, A., Bruegger, E. and Loertscher, S. (2008), 'On cheating, doping and whistleblowing', *European Journal of Political Economy* **24**(2), 415–436.

Berentsen, A. and Lengwiler, Y. (2004), 'Fraudulent Accounting and Other Doping Games', *Journal of Institutional and Theoretical Economics* **160**, 1–14.

Bernard, A. B. and Busse, M. R. (2004), 'Who wins the Olympic Games: Economic resources and medal totals', *The Review of Economics and Statistics* **86**(February), 413–417.

Bird, E. J. and Wagner, G. G. (1997), 'Sport as a Common Property Resource: A Solution to the Dilemmas of Doping', *Journal of Conflict Resolution* **41**(6), 749–766.

Breivik, G. (1992), 'Doping Games A Game Theoretical Exploration of Doping', *International Review for the Sociology of Sport* **27**(3), 235–253.

Buechel, B., Emrich, E. and Pohlkamp, S. (2014), 'Nobody's Innocent: The Role of Customers in the Doping Dilemma', *Journal of Sports Economics* .

Cisyk, J. and Courty, P. (2015), 'Do Fans Care About Compliance to Doping Regulations in Sports? The Impact of PED Suspension in Baseball', *Journal of Sports Economics* .

Clausen, A. (2013), Moral Hazard with Counterfeit Signals.

Cycling Independent Reform Commission (2015), Report to the President of the Union Cycliste Internationale, Technical report.

Dilger, A., Frick, B. and Tolsdorf, F. (2007), 'Are Athletes Doped? Some Theoretical Arguments and Empirical Evidence', *Contemporary Economic Policy* **25**(4), 604–615.

Dilger, A. and Tolsdorf, F. (2005), Karriereverläufe und Doping von 100 m-Läufern, *in* H. Kossbiel and T. Spengler, eds, 'Modellgestützte Personalentscheidungen 8 Dilger', Rainer Hampp, Mering, Germany, pp. 103–117.

Eber, N. (2002), 'Credibility and Independence of the World Anti-Doping Agency: A Barro-Gordon-Type Approach to Antidoping Policy', *Journal of Sports Economics* **3**(1), 90–96.

Eber, N. (2008), 'The Performance-Enhancing Drug Game Reconsidered: A Fair Play Approach', *Journal of Sports Economics* **9**(3), 318–327.

Emrich, E. and Pierdzioch, C. (2013*a*), 'A Note on the International Coordination of Antidoping Policies', *Journal of Sports Economics* .

Emrich, E. and Pierdzioch, C. (2013*b*), 'Joining the international fight against doping', *Applied Economics Letters* **20**(15), 1379–1382.

Fudenberg, D. and Tirole, J. (1991*a*), *Game Theory*, The MIT Press.

Fudenberg, D. and Tirole, J. (1991*b*), 'Perfect Bayesian Equilibrium and Sequential Equilibrium', *Journal of Economic Theory* **53**(2), 236–260.

Greene, W. H. (2004), 'The Behavior of the Fixed Effects Estimator in Nonlinear Models,', *Econometrics Journal* **7**, 98–119.

Haugen, K. (2004), 'The Performance-Enhancing Drug Game', *Journal of Sports Economics* **5**(1), 67–86.

Haugen, K. K., Nepusz, T. and Petróczi, A. (2013), 'The multi-player performance-enhancing drug game.', *PloS one* **8**(5), e63306.

Hermann, A. and Henneberg, M. (2013), Anti-doping systems in sports are doomed to fail: A probability and cost anaylsis.

Huffington Post (2013), 'Lance Armstrong: 1999 Tour De France Doping Cover-Up Instigated By Ex-UCI President Hein Verbuggen'.

Johnson, D. and Ali, A. (2004), 'A Tale of Two Seasons: Participation and Medal Counts at the Summer and Winter Olympic Games', *Social Science Quarterly* **85**(4).

Katz, M. L. (1991), 'Game-Playing Agents: Unobservable Contracts as Precommitments', *The RAND Journal of Economics* **22**(3), 307–328.

Kirstein, R. (2012), 'Doping, the Inspection Game, and Bayesian Enforcement', *Journal of Sports Economics* .

Kofman, F. and Lawarree, J. (1993), 'Collusion in hierarchical agency', *Econometrica: Journal of the Econometric Society* **61**(3), 629–656.

Kräkel, M. (2007), 'Doping and cheating in contest-like situations', *European Journal of Political Economy* **23**(4), 988–1006.

Laffont, J.-J. and Martimort, D. (1999), 'Separation of regulators against collusive behavior', *The Rand journal of economics* **30**(2), 232–262.

Leeds, E. and Leeds, M. (2012), 'Gold, silver, and bronze: Determining national success in men's and women's summer Olympic events', *Journal of Economics and Statistics* **232**(2), 279–292.

Macur, J. (2011), 'N.F.L. Says Players Union Delaying H.G.H. Testing This Season'.
URL: *http://www.nytimes.com/2011/09/29/sports/football/hgh-testing-is-now-unlikely-in-nfl-this-season.html*

Maennig, W. (2002), 'On the Economics of Doping and Corruption in International Sports', *Journal of Sports Economics* **3**(1), 61–89.

Mitchell, H. and Stewart, M. F. (2004), 'Does Drug-Testing Deter Participation in Athletic Events? EPO and the Sydney Olympics', *Chance* **17**(2), 13–18.

Oransky, I. (2008), 'Why Would an Olympic Shooter Take Propranolol?'.
URL: *http://www.scientificamerican.com/article/olympics-shooter-doping-propranolol/*

Pierdzioch, C. and Emrich, E. (2013), 'A note on corruption and national Olympic success', *Atlantic Economic Journal* **41**(4), 405–411.

Pitsch, W. and Emrich, E. (2012), 'The Frequency of Doping in Elite Sport: Results of a Replication Study', *International Review for the Sociology of Sport* **47**(5), 559–580.

Pound, R. W., McLaren, R. H. and Younger, G. (2015), The Independent Commission Report #1, Technical report.

Rahman, D. (2012), 'But Who Will Monitor the Monitor', *American Economic Review* **102**(6), 2767–2797.

Scully, G. (1995), *The Market Structure of Sports*, University of Chicago Press.

Sterken, E. and Kuper, G. (2003), Olympic Participation and Performance since 1896, *in* G. Papanikos, ed., 'The Economics, the Management and the Marketing of Sport', Athens Institute for Education and Research, Athens, Greece, pp. 111–130.

Stigler, G. (1974), The Optimum Enforcement of Laws, *in* G. Becker and W. Landes, eds, 'Essays in the Economics of Crime and Punishment', pp. 55–67.

Striegel, H., Ulrich, R. and Simon, P. (2010), 'Randomized response estimates for doping and illicit drug use in elite athletes', *Drug and Alcohol Dependence* **106**(2-3), 230–232.

Strulik, H. (2012), 'Riding High: Success in Sports and the Rise of Doping Cultures', *The Scandinavian Journal of Economics* **114**(2), 539–574.

Thompson, H. (2012), 'Performance enhancement: Superhuman athletes', *Nature* **487**(7407), 287–9.

USADA (2012), Report on Proceedings under the World Anti-Doping Code and the USADA Protocol. USADA v Lance Armstrong., Technical report.

Van Reeth, D. (2012), 'TV demand for the tour de france: The importance of stage characteristics versus outcome uncertainty, patriotism, and doping', *International Journal of Sport Finance* **8**(1), 39–60.

WADA (2000), 'Health Medicine and Research Committee: Teleconference Meeting Notes, August 8, 2000', Online.

WADA (2013*a*), 2012 Anti-Doping Testing Figures Report, Technical report.

WADA (2013*b*), Lack of Effectiveness of Testing Programs, Technical report.

WADA (2015*a*), 2014 Anti-Doping Testing Figures Report, Technical report.

WADA (2015*b*), World Anti-Doping Code 2015, Technical report.

Werner Pitsch, E. E., Pitsch, W., Emrich, E., Klein, M. and Werner Pitsch, E. E. (2007), 'Doping in Elite Sports in Germany: Results of a www Survey', *European Journal of Sport and Society* **4**(2), 89 – 102.

A DERIVATIONS FOR CHAPTER 2

A.1 DERIVATION OF THE FIRST ORDER CONDITION

The social welfare function can be rewritten as:

$$W(n) = (1-\theta n)aF(\beta) + (\rho-\theta)LF(\beta) - \rho L + (1-\rho n)a\int_{\beta}^{\bar{b}} bf(b)\mathrm{d}b$$

Letting $\beta^* = (1-\theta n^*)/(1-\rho n^*)$:

$$\begin{aligned}
W'(n^*) &= -\theta aF(\beta^*) + (1-\theta n^*)\frac{(\rho-\theta)}{(1-\rho n^*)^2}af(\beta^*) + \frac{(\rho-\theta)^2}{(1-\rho n^*)^2}Lf(\beta^*) \\
&\quad -\rho a\int_{\beta^*}^{\bar{b}} bf(b)\mathrm{d}b - (1-\rho n^*)\frac{(\rho-\theta)}{(1-\rho n^*)^2}a\beta^* f(\beta^*) = 0 \\
&\Rightarrow \left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 Lf(\beta^*) - \theta aF(\beta^*) - \rho a\int_{\beta^*}^{\bar{b}} bf(b)\mathrm{d}b = 0
\end{aligned}$$

Since:

$$\mathrm{E}\Big[b \mid b \geq \beta(n^*)\Big] = \int_{\beta^*}^{\bar{b}} b\frac{f(b)}{1-F(\beta^*)}\mathrm{d}b$$

The first order condition can be rewritten as:

$$\left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 Lf(\beta^*) - \theta aF(\beta^*) - \rho a\big(1-F(\beta^*)\big)\mathrm{E}\Big[b \mid b \geq \beta(n^*)\Big] = 0$$

A.2 DERIVATION OF THE SECOND ORDER CONDITION

The second order condition is:

$$\frac{\rho-\theta}{(1-\rho n^*)^2} f(\beta^*) \times \\ \left[\frac{2\rho(\rho-\theta)}{1-\rho n^*} L + \left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 \frac{f'(\beta^*)}{f(\beta)} L + \rho\beta^* a - \theta a\right] \leq 0$$

This expression can be simplified since:

$$\rho\beta^* a - \theta a = \frac{\rho-\theta}{1-\rho n^*} a$$

Substituting this into the second order condition:

$$\frac{(\rho-\theta)^2}{(1-\rho n^*)^3} f(\beta^*) \left[2\rho L + \frac{\rho-\theta}{1-\rho n^*} \frac{f'(\beta^*)}{f(\beta)} L + a\right] \leq 0$$

Since $\frac{(\rho-\theta)^2}{(1-\rho n^*)^3} f(\beta^*) > 0$, a necessary condition for the second order condition to hold is $f'(\beta^*) < 0$.

A.3 DERIVATION OF COMPARATIVE STATICS

The derivative of the first order condition with respect to μ is:

$$\begin{aligned} \frac{\partial W'(n)}{\partial \mu} &= \left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 \frac{L}{\sigma^2} \alpha^* \phi(\alpha^*) + \frac{\theta a}{\sigma} \phi(\alpha^*) \\ &\quad -\rho a\left[1-\Phi(\alpha^*)\right] - \frac{\rho\mu a}{\sigma} \phi(\alpha^*) - \rho a \alpha^* \phi(\alpha^*) \end{aligned}$$

Using the fact that:

$$\rho\beta^* - \theta = \frac{\rho-\theta}{1-\rho n^*},$$

this can be simplified to:

$$\frac{\partial W'(n)}{\partial \mu} = \left(\frac{\rho-\theta}{1-\rho n^*}\right)^2 \frac{\alpha^*}{\sigma^2} L\phi(\alpha^*) - \frac{\rho-\theta}{\sigma(1-\rho n^*)} a\phi(\alpha^*) - \rho a\left[1-\Phi(\alpha^*)\right]$$

The derivative of the first order condition with respect to σ is:

$$\begin{aligned}\frac{\partial W'(n)}{\partial \sigma} &= \left(\frac{\rho-\theta}{(1-\rho n^*)\sigma}\right)^2 \left(\alpha^2-1\right) L\phi(\alpha^*) + \theta a \frac{\alpha^*}{\sigma}\phi(\alpha^*) \\ &\quad -\rho a \phi(\alpha^*) + \rho a \alpha^* \phi'(\alpha^*) - \rho a \mu \frac{\alpha^*}{\sigma}\phi(\alpha^*)\end{aligned}$$

This can be simplified since $\rho\beta^* - \theta = (\rho-\theta)/(1-\rho n)$ and $\phi'(\alpha) = -\alpha\phi(\alpha)$:

$$\frac{\partial W'(n)}{\partial \sigma} = \left(\frac{\rho-\theta}{(1-\rho n^*)\sigma}\right)^2 \left(\alpha^2-1\right) L\phi(\alpha^*) - \left(\frac{\rho-\theta}{1-\rho n^*}\right) \alpha^* \frac{a}{\sigma}\phi(\alpha^*) - \rho a \phi(\alpha^*)$$

B PROOFS FOR CHAPTER 3

B.1 DERIVATION OF CONSUMERS' EXPECTED UTILITY FUNCTION

Consumers derive utility P from participating where $P \in \{\underline{P}_L, \underline{P}_H, \overline{P}_H\}$ is the performance of the athlete. If the athlete is found to have cheated, consumers lose $R > P \; \forall P$. Therefore, if the athlete tests positive in the ADA's test, consumers do not wish to participate. However, even if the athlete does not test positive in the ADA's test, there is still the possibility that the athlete is found to have cheated in the future, for example due to whistle-blowers. This was true in the case of Lance Armstrong who never publicly failed a drug test.

If an athlete cheated, the probability that he is found to have cheated in the future is α, independent of his performance and the probability with which he was tested. An athlete who did not cheat cannot be found to have cheated in the future. The discounted loss for consumers if cheating is uncovered in the future is δR. Therefore if the athlete cheats and does not test positive, the expected utility of consumers if they participate is:

$$P - \alpha\delta R$$

Therefore if consumers believe the athlete cheated with probability γ, observe performance P and that the athlete did not test positive, their expected utility is:

$$P - \alpha\delta\gamma R$$

They are therefore prepared to participate if:

$$\gamma \leq \frac{P}{\alpha\delta R}$$

It is assumed that $P < \alpha\delta R \; \forall P$ so that consumers will never participate if they know with certainty that the athlete cheated. In addition, since P, α and R are all strictly positive, if consumers know with certainty that the athlete did not cheat, they will participate for all possible performances.

After observing $\underline{P}_H$ and no positive test, consumers participate if:

$$\gamma \leq \frac{\underline{P}_H}{\alpha\delta R} \equiv \hat{\gamma}$$

The final assumption is that $\hat{\gamma} < 1 - \mu$ so that if consumers cannot update their prior after observing $\underline{P}_H$, they do not wish to participate.

The expected utility function of consumers can therefore be summarised as $f(P, \gamma)$ where $f(P, 0) > 0$, $f(\underline{P}_H, 1 - \mu) < 0$ and $f(\underline{P}_H, \hat{\gamma}) = 0$.

B.2 PROOF OF PROPOSITION 3.1

Consumers' behaviour when $\underline{P}_L$, $\overline{P}_H$ or when the athlete tests positive is clear. However, consumers' behaviour after observing $\underline{P}_H$ and that the athlete did not test positive is not deterministic.

Simultaneous timing

In any no-cheating equilibrium it must be the case that neither type of athlete cheats and therefore consumers participate after observing $\{\underline{P}_H, \text{no positive test}\}$. Type L will cheat unless $t \geq t_1$. Therefore a no-cheating equilibrium exists if it is optimal for the ADA to set $t = t_1$. However, when timing is simultaneous and t is not observable, the ADA takes the athlete's decision as given and knows its choice of t cannot directly affect consumers' decision regarding whether or not to participate. Therefore it always benefits from deviating and instead setting $t = 0$ since testing is costly.

Sequential timing

When timing is sequential, the ADA takes into account how its choice of t affects the athlete's decision regarding whether or not to cheat. The only potential equilibrium with no-cheating involves the following strategies where the strategies which would have to be followed on the equilibrium path are in bold:

- ADA: $\boldsymbol{t = t_1}$
- Type H: **Don't cheat**
- Type L: **Don't cheat if $\boldsymbol{t \geq t_1}$**; cheat if $t < t_1$
- Consumers: **Participate if $\underline{P}_L$ or $\underline{P}_H$ and athlete does not test positive**; do not participate if $\overline{P}_H$ or athlete tests positive
- $\gamma = 0$ if $\underline{P}_L$ or $\underline{P}_H$ and the athlete does not test positive; $\gamma = 1$ if $\overline{P}_H$ or the athlete tests positive

The ADA's payoff in this potential equilibrium is:

$$(1 - \mu)c_L + \mu(c_L + d) - t_1 K$$

If neither type of athlete cheats, then the only equilibrium strategy for consumers is to participate after observing $\underline{P}_H$ since they are then certain the athlete is type H. If consumers participate then type L will cheat unless $t \geq t_1$. There is consequently a no-cheating equilibrium if it is optimal for the ADA to set $t = t_1$. It will never be optimal for the ADA to set $t > t_1$. However, if the ADA deviates and sets $t = 0$, type L will then cheat while type H will still not cheat. Consumers cannot observe this deviation and so according to their strategy participate. The ADA's payoff is then:

$$c_L + d$$

This exceeds the ADA's payoff when $t = t_1$. Therefore, since a no-cheating equilibrium could only exist if it was optimal for the ADA to set $t = t_1$, a no-cheating equilibrium does not exist.

B.3 PROOF OF PROPOSITION 3.3

By definition, $\hat{\gamma}$ is the posterior probability which makes consumers indifferent between participating and not. If a type L athlete cheats with probability $\kappa(t)$, after observing $\{t,\ \underline{P}_H,$ no positive$\}$ consumers are indifferent between participating and not if:

$$\begin{aligned}\hat{\gamma} &= \frac{\kappa(t)(1-\mu)(1-\rho t)}{\kappa(t)(1-\mu)(1-\rho t)+\mu} \\ \Rightarrow \kappa(t) &= \frac{\mu\hat{\gamma}}{(1-\hat{\gamma})(1-\mu)(1-\rho t)}\end{aligned}$$

When $t = \hat{t}$, consumers are willing to participate with certainty even if type L cheats with certainty. Therefore if $\hat{t} \leq t < t_1$, $\kappa(t) = 1$. If type L cheats with probability $\kappa(t)$, consumers are indifferent between participating and not for $t \in \left[0, min\{\hat{t}, t_1\}\right)$. When $t \in [0, \hat{t})$, $\kappa(t)$ is strictly increasing in $\hat{\gamma}$ and μ. When $\hat{\gamma}$ increases, consumers are prepared to participate when there is a greater probability that the athlete cheated. Therefore when $\hat{\gamma}$ increases, type L can cheat with a higher probability and consumers are still prepared to mix between participating and not. When μ increases, the prior probability that the athlete is type L declines. Therefore consumers are willing to participate even when type L cheats with higher probability since this is offset by the lower probability of the athlete being type L.

If consumers participate with probability $\sigma(t)$ when $t \in \left[0, min\{\hat{t}, t_1\}\right)$ and the athlete does not test positive, type L is indifferent between cheating and not if:

$$\sigma(t) = \frac{a_L + \rho t F}{(1-\rho t)(a_L + b)}$$

When $t \geq \hat{t}$, consumers participate with certainty even if type L cheats. Therefore $\sigma(\hat{t}) = 1$.

For $t \in [0, \hat{t})$, $\sigma(t)$ is decreasing in b and increasing in F provided $t > 0$. When the punishment for the athlete if he is caught cheating increases, the athlete is only willing to cheat if consumers participate with a higher probability when he does so. Likewise, if b, the athlete's gain from cheating, increases, he is willing to cheat even if it is less likely that consumers will participate.

In any no-cheating equilibrium, consumers must participate with certainty. A no-cheating equilibrium therefore only exists if the ADA prefers to set $t = t_1$ rather than any other value of t. It will never be optimal for the ADA to set $t > t_1$. Likewise, if $\hat{t} < t_1$, it will never be optimal for the ADA to set $t \in (\hat{t}, t_1)$. The ADA therefore chooses between setting t_1 and the optimal t from the range $\left[0, min\{\hat{t}, t_1 - \epsilon\}\right]$. If the ADA sets $t = t_1$, it obtains payoff:

$$(1-\mu)c_L + \mu(c_L + d) - t_1 K$$

If it instead sets t in the range $\left[0, min\{\hat{t}, t_1 - \epsilon\}\right]$ its payoff is:

$$(1-\mu)\left[\Big(1-\kappa(t)\Big)c_L + \kappa(t)\Big(\sigma(t)(1-\rho t)(c_L+d) - \rho t B\Big)\right] + \mu\sigma(t)(c_L+d) - tK$$

The ADA's payoff from setting $t = t_1$ and having neither type of athlete cheat with certainty is greater than its payoff from setting $\tilde{t} < t_1$ and having type L cheat with positive probability if:

$$(1-\mu)\kappa(\tilde{t})\left[c_L + \rho\tilde{t}B - \sigma(\tilde{t})(1-\rho\tilde{t})(c_L+d)\right] + \mu\Big(1-\sigma(\tilde{t})\Big)(c_L+d) - (t_1-\tilde{t})K \ \geq \ 0$$

for all $\tilde{t} \in \left[0, min\{\hat{t}, t_1 - \epsilon\}\right]$.

B.3.1 *Proportional Payoffs*

When $c_L = \eta a_L$, $d = \eta b$ and $B = \eta F$ where $\eta > 0$ and $min\{\hat{t}, t_1\} = t_1$, as $K \to 0$, the ADA's profit from deviating to $t < t_1$ is:

$$(1-\mu)c_L + \mu\frac{c_L + \rho t B}{1-\rho t}$$

which is increasing in t regardless of B. Therefore if the ADA chooses to deviate, it does so by setting $t_1 - \epsilon$. The ADA then prefers not to deviate if:

$$\begin{aligned}(1-\mu)c_L + \mu\frac{c_L + \rho(t_1-\epsilon)B}{1-\rho(t_1-\epsilon)} &\leq (1-\mu)c_L + \mu(c_L+d)\\ \Rightarrow t_1 - \epsilon &\leq \frac{d}{\rho(c_L+d+B)} = \frac{b}{\rho(a_L+b+F)} \equiv t_1\end{aligned}$$

so when $min\{\hat{t}, t_1\} = t_1$, the ADA will never wish to deviate.

If instead $min\{\hat{t}, t_1\} = \hat{t}$ and the ADA sets $t = \hat{t}$, the profit function cannot be simplified in this manner because $\sigma(\hat{t}) \neq (a_L + \rho\hat{t}F)/(1-\rho\hat{t})(a_L+b)$ but instead is equal to 1. Therefore if $min\{\hat{t}, t_1\} = \hat{t}$, as $K \to 0$, the ADA's profit from setting $\hat{t}$ is:

$$(1-\mu)\left[(1-\rho\hat{t})(c_L+d) - \rho\hat{t}B\right] + \mu(c_L+d)$$

The ADA prefers to set $\hat{t}$ rather than $\hat{t} - \epsilon$ if:

$$\begin{aligned}(1-\mu)\left[(1-\rho\hat{t})(c_L+d) - \rho\hat{t}B\right] + \mu(c_L+d)&\\ -(1-\mu)c_L - \mu\sigma(\hat{t}-\epsilon)(c_L+d) &\geq 0\\ \Rightarrow (1-\mu)\left[\frac{d}{\rho(c_L+d+B)} - \hat{t}\right] + \mu\frac{(c_L+d)}{\rho(c_L+d+B)}\left[1-\sigma(\hat{t}-\epsilon)\right] &\geq 0\\ \Rightarrow (1-\mu)\left[t_1 - \hat{t}\right] + \mu\frac{(c_L+d)}{\rho(c_L+d+B)}\left[1-\sigma(\hat{t}-\epsilon)\right] &\geq 0\end{aligned}$$

which is true since $t_1 > \hat{t}$ and $\sigma(\hat{t}-\epsilon) < 1$. Therefore, since it has already been shown that the profit function is increasing in t for $t \in [0, \hat{t}-\epsilon]$, when $\hat{t} < t_1$, if the ADA decides to deviate from setting t_1, it will do so by setting $\hat{t}$.

The ADA would then prefer not to deviate from t_1 to $\hat{t}$ if:

$$\begin{aligned}(1-\mu)c_L + \mu(c_L+d) &\geq (1-\mu)(1-\rho\hat{t})(c_L+d) - \rho\hat{t}(1-\mu)B + \mu(c_L+d)\\ \Rightarrow \hat{t} &\geq \frac{d}{\rho(c_L+d+B)} \equiv t_1\end{aligned}$$

which is not satisfied when $min\{\hat{t}, t_1\} = \hat{t}$. Therefore, when $min\{\hat{t}, t_1\} = \hat{t}$ and the ADA's and athlete's payoffs are proportional, a no-cheating equilibrium does not exist.

However, a no-cheating equilibrium still exists provided B is sufficiently large since as B increases, t_1 decreases and there are fewer values of $\hat{t}$ which are less than t_1.

B.4 PROOF OF PROPOSITION 3.5

The strategies of the players are as follows. Equilibrium strategies are highlighted in bold.

- ADA: $\boldsymbol{t = t_1}$
- Type H athlete: **Don't cheat**
- Type L athlete: **Don't cheat if $\boldsymbol{t \geq t_1}$;** cheat with probability κ_{NT} if $t < t_1$
- Consumers: **Participate if $\boldsymbol{\underline{P}_L}$; participate after $\boldsymbol{\{t \geq t_1,\ \underline{P}_H,}$ negative test$\boldsymbol{\}}$ and $\boldsymbol{\{t \geq t_1,\ \underline{P}_H,}$ no test$\boldsymbol{\}}$;** participate after observing $\{t < t_1,\ \underline{P}_H,$ negative test$\}$; participate with probability $\sigma_{NT}(t)$ after observing $\{t < t_1,\ \underline{P}_H,$ no test$\}$; don't participate after $\overline{P}_H$ or a positive test.
- Beliefs: **$\boldsymbol{\gamma = 0}$ if $\boldsymbol{\underline{P}_L}$ regardless of testing; $\boldsymbol{\gamma = 0}$ if $\boldsymbol{\{t \geq t_1,\ \underline{P}_H,}$ negative test$\boldsymbol{\}}$ or $\boldsymbol{\{t \geq t_1,\ \underline{P}_H,}$ no test$\boldsymbol{\}}$;** $\gamma = \hat{\gamma}$ if $\{t < t_1,\ \underline{P}_H,$ no test$\}$; $\gamma = \frac{(1-\rho)\hat{\gamma}}{1-\rho\hat{\gamma}}$ if $\{t < t_1,\ \underline{P}_H,$ negative test$\}$; $\gamma = 1$ after $\overline{P}_H$ or positive test

Consumers never participate after observing the information set $\{t,\ \underline{P}_H,$ positive test$\}$ since they then know with certainty that the athlete was type L and cheated. Allowing for pure and mixed strategies, there are nine combinations of actions which consumers could take after observing $\{t < t_1,\ \underline{P}_H,$ negative test$\}$ or $\{t < t_1,\ \underline{P}_H,$ no test$\}$. These are summarised in table B.1 and are discussed in turn.

{DON'T PARTICIPATE, DON'T PARTICIPATE} The first entry refers to the action after $\{t < t_1, \underline{P}_H,$ negative test$\}$ and the second to the action after $\{t < t_1,\ \underline{P}_H,$ no test$\}$. For all possible values of t, if consumers do not participate after observing either of these information sets, type L will not cheat. However, consumers then wish to participate after observing either of the information sets since they know the athlete was type H.

	$\{t < t_1,\ \underline{P}_H,$ negative test$\}$	$\{t < t_1,\ \underline{P}_H,$ no test$\}$
1	Don't Participate	Don't Participate
2	Participate	Participate
3	Don't Participate	Participate
4	Mix	Participate
5	Don't Participate	Mix
6	Mix	Mix
7	Participate	Don't Participate
8	Mix	Don't Participate
9	Participate	Mix

Table B.1: Possible combinations of actions taken by consumers

{PARTICIPATE, PARTICIPATE} When $t \geq t_1$, even when consumers behave in this manner type L will not choose to cheat. Strategies and beliefs are then consistent. However, if $t < t_1$ type L will cheat given this behaviour by consumers. Consumers then definitely do not want to participate after observing $\{t < t_1,\ \underline{P}_H,$ no test$\}$ and may or may not wish to participate after observing $\{t < t_1,\ \underline{P}_H,$ negative test$\}$.

{DON'T PARTICIPATE, PARTICIPATE}, {MIX, PARTICIPATE}, {DON'T PARTICIPATE, MIX} AND {MIX, MIX} After observing $\{t < t_1,\ \underline{P}_H,$ negative test$\}$ consumers believe the athlete is type L with probability:

$$\gamma = \frac{(1-\mu)(1-\rho)\kappa(t)}{(1-\mu)(1-\rho)\kappa(t)+\mu} \equiv \gamma_1$$

where $\kappa(t)$ is the probability with which type L cheats.

After observing $\{t < t_1,\ \underline{P}_H,$ no test$\}$ their belief is:

$$\gamma = \frac{(1-\mu)\kappa(t)}{(1-\mu)\kappa(t)+\mu} \equiv \gamma_2$$

When making his decision regarding whether or not to cheat, the athlete does not know whether or not he will be tested. Therefore his strategy cannot be contingent on this, although it can be contingent on the probability with which he will be tested. Consequently $\kappa(t)$ is the same in γ_1 and γ_2 and since $\rho > 0$, $\gamma_1 < \gamma_2$ provided $\kappa(t) > 0$. Consumers participate with positive probability if $\gamma \leq \hat{\gamma}$. If consumers participate with positive probability after observing no test, it then must be the case that they wish to participate with certainty after observing a negative test since $\gamma_2 \leq \hat{\gamma} \Rightarrow \gamma_1 < \hat{\gamma}$. Intuitively, if consumers were sufficiently confident to participate without observing a test result, they must be even more confident if they also observe a negative test. Consequently {Don't Participate, Participate}, {Mix, Participate}, {Don't Participate, Mix} and {Mix, Mix} are not consistent for any values of t.

{PARTICIPATE, DON'T PARTICIPATE} AND {MIX, DON'T PARTICIPATE}

Consumers wish to participate with positive probability after observing $\{t < t_1,\ \underline{P}_H,$ negative test$\}$ but not after observing $\{t < t_1,\ \underline{P}_H,$ no test$\}$ if:

$$\frac{(1-\mu)(1-\rho)\kappa(t)}{(1-\mu)(1-\rho)\kappa(t)+\mu} \leq \hat{\gamma} < \frac{(1-\mu)\kappa(t)}{(1-\mu)\kappa(t)+\mu}$$

where $\kappa(t)$ is the probability with which type L cheats. This condition can only be satisfied if $\kappa(t) > 0$ so type L must cheat with strictly positive probability.

Given this behaviour by consumers, for all $t \leq 1$, if the athlete prefers not to cheat when consumers participate provided he does not test positive, he definitely prefers not to cheat when consumers only participate with positive probability if he tests negative. Therefore, for all $t \geq t_1$, type L does not cheat with certainty if consumers behave in this manner and consequently this strategy is not consistent.

If consumers participate with positive probability after $\{t < t_1,\ \underline{P}_H,$ negative test$\}$ but do not participate after observing $\{t < t_1,\ \underline{P}_H,$ no test$\}$, type L prefers to cheat with strictly positive probability if:

$$\begin{aligned} t\sigma_{-ve}(t)(1-\rho)(a_L+b) - \rho t F &\geq a_L \\ \sigma_{-ve}(t) &\geq \frac{a_L + \rho t F}{(1-\rho)(a_L+b)t} \equiv \sigma_1 \end{aligned}$$

$\sigma_{-ve}(t)$ is the probability with which consumers participate after observing a negative test.

Therefore type L prefers to cheat with positive probability when consumers behave in this manner if $\sigma_{-ve}(t) \geq \sigma_1$. Such a σ only exists if $\sigma_1 \leq 1$:

$$\begin{aligned} \frac{a_L + \rho t F}{(1-\rho)(a_L+b)t} &\leq 1 \\ \Rightarrow \big[(1-\rho)(a_L+b) - \rho F\big]t &\geq a_L \end{aligned}$$

If $(1-\rho)(a_L+b) - \rho F < 0$, then this condition can only be satisfied if $t < 0$ which is not possible. If $(1-\rho)(a_L+b) - \rho F > 0$, this condition is satisfied if:

$$t \geq \frac{a_L}{(1-\rho)(a_L+b) - \rho F} \equiv t_2$$

It has been shown above that consumers will not behave in this manner if $t \geq t_1$. Therefore these strategies and beliefs can only be consistent if $t_2 < t_1$. It can be shown that $t_2 < t_1 \Rightarrow t_1 > 1$:

$$\begin{aligned} t_2 &< t_1 \\ \Rightarrow \frac{a_L}{(1-\rho)(a_L+b) - \rho F} &< \frac{b}{\rho(a_L+b+F)} \\ \Rightarrow \rho a_L(a_L+b+F) &< b(a_L+b) - \rho b(a_L+b+F) \\ \Rightarrow \rho(a_L+b)(a_L+b+F) &< b(a_L+b) \\ \Rightarrow \frac{b}{\rho(a_L+b+F)} &> 1 \\ \Rightarrow t_1 &> 1 \end{aligned}$$

However, it is assumed that $t_1 \leq 1$ since when $t_1 > 1$, the ADA cannot deter cheating by both types even if it wanted to and a no-cheating equilibrium will never exist. Therefore, these beliefs and strategies are not consistent when $t_1 \leq 1$.

{PARTICIPATE, MIX} Consumers are indifferent between participating and not participating after observing no test if $\gamma_2(t) = \hat{\gamma}$. This is the case when:

$$\kappa = \frac{\mu\hat{\gamma}}{(1-\hat{\gamma})(1-\mu)} \equiv \kappa_{NT}$$

γ_2 is strictly increasing and continuous in κ. When $\kappa = 1$, $\gamma_2 = 1 - \mu > \hat{\gamma}$. When $\kappa = 0$, $\gamma_2 = 0 < \hat{\gamma}$. Therefore there always exists $\kappa_{NT} \in (0,1)$. If $\gamma_2 = \hat{\gamma}$ then $\gamma_1 < \hat{\gamma}$ and consumers participate with certainty after observing a negative test.

If consumers participate with certainty after a negative test and mix after no test, type L is indifferent between cheating and not if:

$$\begin{aligned} t(1-\rho)(a_L+b) + \sigma_{NT}(1-t)(a_L+b) - \rho t F &= a_L \\ \Rightarrow \sigma_{NT}(t) &= \frac{a_L + \rho t F - t(1-\rho)(a_L+b)}{(1-t)(a_L+b)} \end{aligned}$$

where $\sigma_{NT}(t)$ is the probability with which consumers mix after observing $\{t < t_1,\ \underline{P}_H,\ \text{no test}\}$. This is only possible if $\sigma_{NT}(t) < 1$:

$$\begin{aligned} \frac{a_L + \rho t F - t(1-\rho)(a_L+b)}{(1-t)(a_L+b)} &< 1 \\ \Rightarrow t &< \frac{b}{\rho(a_L+b+F)} \equiv t_1 \end{aligned}$$

Therefore strategies and beliefs are consistent if consumers participate after $\{t < t_1,$ $\underline{P}_H$, negative test$\}$ and mix with probability $\sigma_{NT}(t)$ after $\{t < t_1,\ \underline{P}_H,\ \text{no test}\}$. If the ADA sets $t < t_1$, it then obtains expected payoff:

$$\begin{aligned} &(1-\mu)\Big[\kappa_{NT}\Big(t(1-\rho)(c_L+d) + \sigma_{NT}(t)(1-t)(c_L+d) - \rho t B\Big) + (1-\kappa_{NT})c_L\Big] \\ &+\mu\Big[t + \sigma_{NT}(t)(1-t)\Big](c_L+d) - tK \end{aligned}$$

This is less than the ADA's profit from setting $t = t_1$ and having neither type of athlete cheat with certainty if:

$$\begin{aligned}&(1-\mu)c_L + \mu(c_L + d) - t_1 K \geq \\ &(1-\mu)\Big[\kappa_{NT}\Big(\tilde{t}(1-\rho)(c_L+d) + \sigma_{NT}(\tilde{t})(1-\tilde{t})(c_L+d) - \rho\tilde{t}B\Big) + (1-\kappa_{NT})c_L\Big] \\ &+\mu\Big[\tilde{t} + \sigma_{NT}(\tilde{t})(1-\tilde{t})\Big](c_L+d) - \tilde{t}K\end{aligned}$$

for all $\tilde{t} \in [0, t_1)$. This can be simplified to:

$$\begin{aligned}(1-\mu)\kappa_{NT}\Big[c_L + \rho\tilde{t}B - \Big(\sigma_{NT}(\tilde{t})(1-\tilde{t}) + (1-\rho)\tilde{t}\Big)(c_L+d)\Big]& \\ +\mu(1-\tilde{t})\Big(1-\sigma_{NT}(\tilde{t})\Big)(c_L+d) - (t_1-\tilde{t})K &\geq 0\end{aligned}$$

which is condition 3.2 in proposition 3.5.

B.4.1 *Proportional Payoffs*

Again, it is possible to conduct further analysis if the simplifying assumption is made that $c_L = \eta a_L$, $d = \eta b$ and $B = \eta F$. As $K \to 0$, the ADA's profit if it sets t in the range $[0, t_1)$ is then:

$$c_L + \mu\rho t(c_L + d + B)$$

which is again increasing in t. The ADA will therefore set $t = t_1 - \epsilon$ if it decides to deviate. The ADA then prefers not to deviate if:

$$\begin{aligned}c_L + \mu\rho(t_1-\epsilon)(c_L+d+B) &\leq (1-\mu)c_L + \mu(c_L+d) \\ \Rightarrow t_1 - \epsilon &\leq \frac{d}{\rho(c_L+d+B)} = \frac{b}{\rho(a_L+b+F)} \equiv t_1\end{aligned}$$

which is always satisfied.

B.5 PROOF OF PROPOSITION 3.6

If the ADA sets $t = t_1$, its payoff is:

$$\Pi(t = t_1) = (1 - \mu)c_L + \mu(c_L + d) - t_1 K$$

However, when timing is simultaneous neither consumers nor the athlete can respond to a deviation by the ADA. Therefore by setting $t = 0$ the ADA expects to obtain profit:

$$\Pi(t = 0) = (1 - \mu)c_L + \mu(c_L + d) > \Pi(t = t_1)$$

B.6 PROOF OF PROPOSITION 3.7

A no-cheating equilibrium exists provided it is optimal for the ADA to set $t = t_1$. Once again, if the ADA sets $t = t_1$, its payoff is:

$$\Pi(t = t_1) = (1 - \mu)c_L + \mu(c_L + d) - t_1 K$$

The ADA recognises that if it deviates to $\tilde{t} < t_1$, type L will cheat. However, consumers cannot observe this deviation and so participate provided type L does not test positive. As $K \to 0$, the ADA's payoff from setting $\tilde{t} < t_1$ is then:

$$\Pi(t = \tilde{t}) = (1 - \mu)\Big[(1 - \rho\tilde{t})(c_L + d) - \rho\tilde{t}B - \rho\tau(1 - \rho\tilde{t})D\Big] + \mu(c_L + d)$$

This profit function is linear in $\tilde{t}$. Therefore given $t < \tilde{t}$, the ADA will either wish to set $\tilde{t} = 0$ or $\tilde{t} = t_1 - \epsilon$. If $\rho\tau D < c_L + d + B$ the profit function for $\tilde{t} < t_1$ is decreasing in $\tilde{t}$ and if the ADA wishes to deviate, it will do so by setting $\tilde{t} = 0$. If instead $\rho\tau D \geq c_L + d + B$ the profit function is increasing in $\tilde{t}$ and the ADA prefers to set $\tilde{t} = t_1 - \epsilon$ if it decides to deviate.

The ADA prefers to set $t = t_1$ rather than any $\tilde{t} < t_1$ if:

$$\tau \geq \frac{1}{\rho D}\left[d - \frac{\rho\tilde{t}(c_L + B)}{1 - \rho\tilde{t}}\right] \equiv \underline{\tau}(\tilde{t})$$

When $\rho\tau D < c_L + d + B$, $\tilde{t} = 0$ and the condition for a no-doping equilibrium is then:

$$\tau \geq \frac{d}{\rho D} \equiv \underline{\tau}(0)$$

When $\rho\tau D \geq c_L + d + B$, $\tilde{t} = t_1 - \epsilon$ and the condition for a no-doping equilibrium is:

$$\tau \geq \frac{1}{\rho D}\left[d - \frac{\rho(t_1 - \epsilon)(c_L + B)}{1 - \rho(t_1 - \epsilon)}\right] \equiv \underline{\tau}(t_1 - \epsilon)$$

Both $\underline{\tau}(0)$ and $\underline{\tau}(t_1 - \epsilon)$ are decreasing in D. Therefore, provided D is sufficiently large, both $\underline{\tau}(0)$ and $\underline{\tau}(t_1 - \epsilon)$ will be less than 1 and therefore WADA will always be able to implement a sufficiently high τ that the ADA does not wish to deviate.

C APPENDIX FOR CHAPTER 4

C.1 COEFFICIENTS ON PROPORTION OF MEDALS WON IN 2008

	I	II	III	IV
Method	OLS	GLM	OLS	GLM
Observations	All	All	SPA <53	SPA <53
Aquatics	0.795***	13.12***	0.824***	13.21***
	(0.142)	(1.175)	(0.141)	(1.247)
Archery	0.697***	9.516***	0.696***	9.538***
	(0.0744)	(1.711)	(0.0755)	(1.685)
Athletics	0.841***	19.19***	0.860***	19.32***
	(0.144)	(3.087)	(0.143)	(3.126)
Badminton	0.806***	7.276***	0.804***	7.455***
	(0.175)	(0.408)	(0.175)	(0.491)
Boxing	0.595***	15.59***	0.595***	16.25***
	(0.137)	(3.673)	(0.145)	(4.423)
Canoe	0.748***	17.63***	0.743***	17.43***
	(0.134)	(4.052)	(0.133)	(4.224)
Cycling	0.557***	7.035***	0.562***	6.791***
	(0.0763)	(0.900)	(0.0743)	(0.900)
Equestrian	0.640***	8.820***	0.633***	8.716***

	(0.229)	(2.536)	(0.226)	(2.452)
Fencing	0.535**	8.639***	0.530**	8.757***
	(0.259)	(2.548)	(0.259)	(2.552)
Gymnastics	0.563***	7.643***	0.559***	7.901***
	(0.101)	(1.547)	(0.0987)	(1.677)
Judo	0.593***	16.44***	0.584***	18.97***
	(0.124)	(3.403)	(0.133)	(4.326)
Other	0.546***	8.969***	0.545***	8.925***
	(0.112)	(1.179)	(0.120)	(1.285)
Rowing	0.818***	15.42***	0.858***	15.58***
	(0.165)	(3.003)	(0.163)	(3.206)
Sailing	0.548***	10.23***	0.549***	10.33***
	(0.107)	(2.202)	(0.108)	(2.294)
Shooting	0.571***	14.88***	0.563***	15.24***
	(0.148)	(2.240)	(0.152)	(2.339)
Table Tennis	0.753***	6.622***	0.749***	6.701***
	(0.0295)	(0.920)	(0.0296)	(0.938)
Taekwondo	0.408***	11.16***	0.424***	12.50***
	(0.0738)	(3.211)	(0.0755)	(3.815)
Tennis	0.462*	7.182**	0.473*	7.719**
	(0.250)	(3.177)	(0.253)	(3.504)
Triathlon	0.155	3.973	0.196	3.629
	(0.179)	(3.038)	(0.159)	(2.246)
Volleyball	0.628***	8.355***	0.626***	8.296***
	(0.199)	(1.213)	(0.199)	(1.201)
Weightlifting	0.706***	25.53***	0.676***	27.55***

	(0.0742)	(5.457)	(0.190)	(7.002)
Wrestling	0.784***	23.13***	0.820***	26.29***
	(0.123)	(5.123)	(0.123)	(6.501)
Observations	1673	1673	1634	1634

Table C.1: Coefficients on proportion of medals won in 2008.
Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.
The regression also included samples per athlete in 2012 and country fixed effects.