CSEF
Centre for Studies in Economics and Finance

WORKING PAPER NO. 196

Judicial Errors, Legal Standards and Innovative Activity

Giovanni Immordino and Michele Polo

February 2008

This version June 2010

University of Naples Federico II | University of Salerno | Bocconi University, Milan

CSEF - Centre for Studies in Economics and Finance
DEPARTMENT OF ECONOMICS – UNIVERSITY OF NAPLES
80126 NAPLES - ITALY
Tel. and fax +39 081 675372 – e-mail: csef@unisa.it

WORKING PAPER NO. 196

Judicial Errors, Legal Standards and Innovative Activity

Giovanni Immordino[♥] **and Michele Polo**[♠]

Abstract

We analyze the effects of judicial errors on the innovative activity of firms. Successful research investment allows to take a new action that may be ex-post welfare enhancing or welfare decreasing (illegal). Deterrence in this setting works by affecting both the incentives to invest in research and the ex-post choice of the action (marginal deterrence). The two goals may contrast each other and their relative importance shapes the optimal policy. For increasing probabilities of social harm, the enforcer initially promotes research and disregards marginal deterrence: in this case no accuracy is chosen and the policy adopts first a per-se legality and then a (moderately enforced) per-se illegality rule. Conversely, for higher likelihood of social harm, the enforcer favors marginal deterrence eliciting the ex-post efficient actions at the cost of depressing firm's investment: the optimal policy calls for positive .nes and improved accuracy, with a stronger effort to reduce type-I than type-II errors and the adoption of asymmetric protocols of investigation. Improved accuracy allows to discriminate between beneficial and harmful actions, an instance of effect-based legal standard.

Keywords: norm design, innovative activity, enforcement, errors.

JEL classification: D73, K21, K42, L51.

Acknowledgements: We are indebted to Nuno Garoupa, Yannis Katsoulakos, Dilip Mookherjee, Massimo Motta, Marco Pagano, Patrick Rey, Matteo Rizzolli, Lars- Hendrik Röller, Maarten Schinkel and Giancarlo Spagnolo for helpful discussions. Thanks also to seminar participants in Bologna, Bonn, Naples, Milan Bocconi, Venice, Milan Bic- occa (Intertic Conference), the 2007 Italian Congress in Law and Economics (Milan), the 2008 CSEF-IGIER Symposium on Economics and Institutions (Capri), the 2008 ESEM meeting (Milan), the 2008 EARIE meeting (Toulouse) and the 2008 WZB Conference on Deterrence in Competition Policy (Berlin). All usual disclaimers apply.

♥ Università di Salerno and CSEF – Corresponding address: Università di Salerno, 84084 Fisciano (SA), Italy, giimmo@tin.it

♠ Università Bocconi, IGIER and CSEF – Corresponding address: Università Bocconi, Via Sarfatti 25, 20136 Milan, Italy, michele.polo@unibocconi.it

Table of contents

1. Introduction

2. The model

3 . First best

4. Firm's choices: actions and research

5. Optimal enforcement policies

5.1. The status quo policy (SQ)

5.2. The e¢ cient action policy (EA)

5.3. The new action policy (NA)

5.4. Optimal policy with respect to β

5. Conclusions

Appendix

References

1 Introduction

The purpose of this paper is to study the effects of judicial errors on firms' innovative activity. Quite often norms are required to rule delicate issues involving innovative environments, such as the design of liability rules for genetically modified (GM) organisms or the application of antitrust norms in high-tech industries. In these settings, there is widespread concern about the long term impacts of the design and enforcement of norms on the rate of innovative activity. For instance, a recurrent theme in competition policy claims that antitrust legislation should prevent abuses by dominant firms without chilling competition on the merits. Moreover, in enforcing rules applied to novel and innovative issues, errors becomes more likely than in standardized situations, and judicial errors are a major concern.

In order to address these issues, we argue that an enrichment of the traditional model of law enforcement is needed, applying then this richer set-up to judicial errors. Let us discuss the two steps in a sequence.

Law enforcement in innovative environments. The traditional approach of the Law and Economics (L&E) literature is not appropriate to address these problems because, in its basic set-up, it considers the choices of private agents among a set of feasible actions, some of which may cause a social damage and are therefore considered unlawful. In this setting the feasible actions are perfectly known and implementable by the individuals, the only restraint from taking harmful acts being the expected fines associated to illegal practices. The analysis focusses on the ability of law enforcement to influence individuals in their choice among (harmful) actions, which represents the very notion of marginal deterrence[1].

This setting does not allow to fully address the issues we want to analyze. Consider the design and application of legal rules in the above mentioned examples referred to the liability issues in the production of genetically modified seeds, or to the antitrust issues in the design of new products by dominant companies. In both cases the traditional problem, in which the private agents choose among a set of known actions, corresponds to the final stage of a process that initially requires investment in research to identify innovative solutions among which the choice will be made in the end.[2], [3] When firms commit to a research investment,

[1]See Stigler (1970), Shavell (1992), Mookherjee and Png (1994), among others.

[2]To our knowledge, Kaplow (1995) is the only other paper where the design of the law affects agents' learning decisions. In his setting, more complex rules allow better control over individual behaviour but are more difficult for people to understand ex ante and for courts to apply ex post. In Kaplow's setting, individuals can choose not to learn, and to take actions ignoring the associated effects (and fines). Our model differs in that new actions can be taken only upon learning.

[3]This setting has some common elements with the so called "activity level" model: see, for instance, Shavell (1980) and (2007) and Shavell and Polinsky (2000). According to this approach, private benefits and social harms depend on two different decisions of private agents: a level of activity (how long the individual

they still ignore its outcome and the effects of the new action that research may discover. Hence, they may be unable to predict if the new actions will be considered lawful, according to the existing norms.

Two features characterize these more complex situations: firstly, private agents have a richer set of decisions, as initially they must choose how much to invest in research and then, if this is successful, select one among the innovative actions available; secondly, during the innovative process the private agents not only discover how to implement the new actions, but they may also learn the true state of the world, that is whether these latter will be considered *ex-post* as lawful or unlawful according to the prescriptions of the legal rules.

The *ex-ante* inability to evaluate with certainty the social consequences (or legality) of the innovative actions may be due to different reasons. Uncertainty may be rooted in the very nature of the research activity that the firm has to perform, so that the features of the innovation are unknown until discovery. For instance, in the example of the biotech firm, experiments with a new GM seed may promise higher yields but may also pose unknown risks to public health, that can be properly verified only once the research project has been concluded. In other instances, uncertainty may derive from the interaction of the innovation, whose properties may have been controlled and planned by the firm with sufficient confidence, with the economic or social environment at the time the innovation is introduced. The future features of this environment, in turn, will depend on the decisions of a very high number of other agents and cannot be assessed *ex-ante* with certainty. In our second example, a dominant software company may invest in research to tie a new software application into a new personal computer operating system: beyond the initial intent of the company, the efficiency and foreclosure effects of this new software will depend, at the time of its commercial introduction, on the alternative packages and applications available from competitors, which may be only imperfectly foreseen at the time of the research investment.

In this class of situations, deterrence works through an additional channel, by affecting the initial incentives to invest in research: if private agents expect a very restrictive legal treatment of the results of their innovative effort, they will have lower incentives to commit resources to research. As a result, the innovative actions will be discovered and possibly chosen with a lower probability.

Immordino, Pagano and Polo (2009) propose an analytical framework to address these two features of deterrence, analyzing the choice among different policy regime, namely

drives) and a level of precaution (the speed at which the agent drives). This literature focusses mainly on comparison of different liability rules (strict vs fault-based). In our paper the innovative effort resembles to the activity while the choice of new actions parallels precaution. However, in our setting, the time and information structure is different, since innovative effort is taken before uncertainty is resolved and before actions are chosen, while in the "activity level" model activity and precaution are chosen together and uncertainty plays no role.

laissez-faire, ex-post law enforcement and ex-ante authorizarion. The interplay of marginal deterrence on the new actions and deterrence on research investment shapes the optimal policies for different priors on the expected harm of innovations. The impact of antitrust enforcement in innovative industries is analyzed also in a paper by Segal and Whinston (2007). Considering a sequence of innovations, the authors analyze the trade-off between protecting the incumbents, by so doing increasing the rents of the winner and the incentives to invest in innovation, and protecting the innovative entrants, that increases the rate of technical progress. They derive conditions under which the latter effect is the dominant one.

Errors. The two previous papers offer interesting results on law enforcement when innovative activity is a crucial component, but they do not consider judicial errors, the second ingredient of this work. Judicial errors and their reduction, i.e. accuracy, are a central concern in law enforcement: they have been analyzed in the standard model of law enforcement proposed by Kaplow (1994), Kaplow and Shavell (1994, 1996), Polinsky and Shavell (2000) and Png (1986) among others, which focusses on the (negative) impact of such errors on marginal deterrence. In this framework, accuracy is always desirable, and it is chosen optimally balancing the marginal benefits and costs.

A more specific literature on competition policy enforcement has considered the effects of an inappropriate intervention by an antitrust authority. In a model of collusion, Schinkel and Tuinstra (2006), find that the incidence of anti-competitive behavior increases in both types of enforcement errors: type II errors reduce expected fines, while type I errors encourage industries to collude when faced with the risk of a false conviction. Therefore, in their framework as in the general literature on law enforcement accuracy is always desirable. In Katsoulakos and Ulph (2009) a welfare analysis of legal standard is developed, comparing per-se rules and discriminating (effect based) rules characterized by a lower probability of errors. The authors identify some key elements that can help deciding the more appropriate legal standard and the cases in which type-I or type-II accuracy is more desidable. In both papers, as in the literature on accuracy quoted above, the general setting is consistent with the standard set-up, in which the enforcer aims at influencing the choice of firms among (known and feasible) actions, while the issue of innovative activity is not considered.

Our contribution. From the previous discussion, we argue that the existing papers on law enforcement in innovative environments do not consider judicial errors while these latter have been studied in a set-up that is not fit to consider innovative environments. The combination of these ingredients is the purpose of this paper. We adopt a set-up similar to Immordino, Pagano and Polo (2009) introducing the possibility of judicial errors.

Errors in our model occur when assessing the social consequences (legality) of the new action, the more compelling task for the enforcer. Following the literature, we can distinguish two types of errors: the enforcer can erroneously fine the firm when taking the new (socially beneficial) action or mistakenly acquit the (socially damaging) new action.[4]

The enforcer sets the fine and determines, through costly effort, the (type-I or type-II) accuracy of the decision. In our setting, when the enforcer does not pursue any accuracy she is (implicitly) adopting a per-se (legality or illegality) rule that does not discriminate the effects, a property that instead can be optained through accuracy implementing an effect-based rule. Hence, we are able not only to derive the optimal use of the policy instruments but also the optimal legal standards.[5]

The main findings of our analysis can be summarized as follows. The enforcer shapes the policy considering two effects: the impact of the fines and accuracies on the firm's choice of actions (marginal deterrence) and on its research investment. Concerning marginal deterrence we can identify three different cases that are implementable through the policy: the enforcer may prevent the choice of the innovative action in any state of the world (status quo), she may implement in each contingency the efficient action (efficient action), or she may allow the firm to choose the new action in any case (new action). We show that the first two cases involve underinvestment in research compared to the first best, while in the third case the research investment may be higher or lower than the welfare maximizing level.

In the new action case the enforcer gives up the possibility of influencing the firm's choice of action and focusses on enhancing the incentives to research investment. This comes out to be the more desirable option when innovation is ex-ante very beneficial. Indeed, when the new action is very likely to produce a social gain, the optimal policy entails zero fines and no accuracy, an example of per-se legality rule (*laissez-faire*). As the probability of social harm (unlawfulness) increases, fines are progressively raised, still not improving accuracy, moving to a moderately enforced per-se illegality rule. At some point, the optimal policy shifts from the new to the efficient action case, and marginal deterrence starts shaping the policy: when the new action is more likely to produce a social damage, it becomes optimal to prevent its adoption when harmful. Now the enforcer has to pursue two goals, sustaining the research investment and preventing the firm from choosing the new action when socially harmful; the optimal policy requires to use both tools – fines and accuracy. However, type-I errors are reduced more than type-II errors, a case of asymmetric accuracy. In this case the

[4] The first case corresponds to a type-I error in statistical inference and is labelled as a case of over-enforcement or false positive in the L&E jargon, while the second entails a type-II error and involves under-enforcement and false negative.

[5] The discussion on legal standards has occupied an important place in the recent debate on the enforcement of competition policy in preventing foreclosure and monopolization, comparing per-se rules (form-based) and discriminating rules (effect-based). See on this point Gual et al (2005) and Katsoulakos and Ulph (2009).

firm is fined with different probability according to the social consequences of the action, an instance of effect-based legal standard. Finally, when the new action is very likely to be harmful, the policy reduces progressively the research investment up to complete deterrence of research, coming back to per-se prohibition.

The remainder of this paper is organized as follows. Section 2 presents the model. Section 3 sets the first best benchmark. Section 4 analyzes the firm's choices regarding the action and the investment in research. Section 5 and its subsections identify the optimal polices. Section 6 concludes. All the proofs are in the Appendix.

2 The model

We consider a model with a profit-maximizing firm and a benevolent enforcer that may commit mistakes. The firm can choose one among several known and lawful actions or invest in learning to identify a new action, whose private and social effects are *ex-ante* unknown.

The key issue that we wish to explore is the optimal design of fines and accuracy when private innovative activity is important and enforcers are subject to judgement errors.

The firm can choose the *status-quo* action a_0 (planting traditional seeds, offering an untied application) with associated profits $\Pi(a_0)$ and welfare $W(a_0)$: we normalize these two measures to zero, i.e. $\Pi(a_0) = W(a_0) = 0$. Action a_0 is the most profitable among the known and legal actions that the firm is able to implement without investing in learning. It is correctly recognized by the enforcer in terms of its nature (a_0) and social consequences ($W(a_0)$).

Alternatively, if the research activity has been successful, the firm can consider a new action a (*innovation*), with associated profit $\Pi(a) = \Pi > 0$.[6] Depending on the state of nature s, the social consequences of the new action differ. With probability β, a bad state $s = b$ occurs, and the new action has a negative social externality, $W(a) = W_b(a) = \underline{W} < 0$. In this case, private incentives conflict with social welfare. With probability $1 - \beta$, instead, a good state $s = g$ materializes and the new action improves welfare, $W(a) = W_g(a) = \overline{W} > \Pi > 0$. In this case, there is no conflict between private and social incentives, since

[6] In this paper we consider just one possible new action resulting from the learning effort, rather than a set of new actions. In Immordino, Pagano and Polo (2009) we analyze, in the case of no errors, the case of multiple actions and the issue of marginal deterrence. Even in the present simplified environment, however, an issue of marginal deterrence arises since the enforcer can influence the choice between the new action a and the status quo action a_0. Notice that the action that is privately more profitable (a) is also the one that is more socially damaging in the bad state. Since errors make the analysis more complex, we prefer to maintain the model simple without losing insights on this important issue of law enforcement.

the innovation improves both the profits of the firm and social welfare. The three elements β, $\underline{W}$ and $\overline{W}$ summarize the "economic model" considered by the firm and the enforcer, that describes the ex-ante possible outcomes of the innovative activity. We assume that the enforcer and the firm share the same priors on these effects. In our examples, β is the prior probability that the GM seeds will be hazardous to public health, or that the new tied software application, when introduced in the market, will foreclose alternative software packages.

While the firm knows from the beginning how to implement the status-quo action a_0, carrying out the new action requires an investment in learning (experiment with GM seeds, create a new tied application), which accordingly will be referred to as "*research*". If the research investment is successful, the firm will discover how to implement the new action a. In this case, the firm also learns the state of nature s, i.e. whether its innovation is socially harmful or beneficial. Proceeding with our examples, the biotech company learns not only how to produce new GM seeds, but also the dangers that they pose to public health and the damages that it might face under the prevailing liability rules. And the software company, once the new application is created, is able to predict whether in the current market conditions it will foreclose the alternative packages or not.

The amount of resources I that the firm invests in research determines its chances of success: for simplicity, the firm's probability $p(I)$ of learning how to carry out the new action a is assumed to be linear in I, i.e. $p(I) = I$ with $I \in [0, 1]$. The cost of learning is increasing and convex in the firm's investment. For simplicity we assume $c(I) = c\frac{I^2}{2}$ with

$$c > \overline{W} \tag{1}$$

to ensure an internal solution in all regimes.

The institutional framework in the design and enforcement of norms is as follows. The legislator writes the legal rule, which specifies under what circumstances the actions are legal or not, and the admitted fine. The enforcement officials seek evidence on the firm's action and on the associated social consequences, and may commit errors in this latter task. Since the legislator and the enforcer are assumed to be benevolent, we can analyze their choice as if they are taken by a single agent, that we call from here on the enforcer.

The legal rule identifies some circumstances that make the new action legal or unlawful. In general we can adapt this setting to a wide range of formal frameworks: for instance, the norm may state that the new action is illegal whenever it occurs together with contingencies $x_1, ..., x_n$, a case that is reminiscent of more or less articulated *per-se* rules.[7] Alternatively, illegality may be related to the effects of the action, as required under a rule of reason

[7] Drawing on antitrust legislation, for instance, a very simple per-se rule would consider as illegal the practice of resale price maintenance when adopted by a firm with a market share larger than x%. A more

approach. It is important to stress that our analysis can be adapted to either of the two cases. All that matters is that, at the time of the innovative investment, the elements that the legal rule identifies in order to assess the lawfulness of the new action are not known with certainty. With this important caveat in mind, we consider a legal rule written as follows, which allows us to greatly simplify the notation in the analysis:

> The action a_0 is lawful; the (new) action a is illegal if *ex-post* socially damaging, i.e. if $s = b$ and therefore $W(a) < 0$. The illegal action is sanctioned according to a fine f chosen in the interval $[0, F]$.

For instance, the legal rule prohibits the commercialization of hazardous seeds or the adoption of practices that foreclose the market to competitors.

The enforcer has first to identify the action chosen (a_0 or a) and the social consequences of the action (0 or $W_s(a)$ for $s = g, b$). Obtaining evidence on these elements requires to commit resources. We define respectively as *monitoring* and *accuracy* the activities devoted to obtaining evidence on the action chosen and on its consequences (legality). In order to focus the analysis on accuracy, we assume for simplicity but without loss of generality that the resources devoted to monitoring are exogenously given and such that the enforcer correctly recognizes with probability 1 the action chosen. When evaluating the consequences (lawfulness) of the action, instead, it might commit errors and may invest in accuracy to make its assessment more precise.[8]

More precisely, we assume that the enforcer is more accurate in assessing the effects of the status-quo rather than the new action, that is judicial errors occur only when assessing the effects (or, equivalently, the state of nature s) of the new action a, while the status-quo action a_0 is correctly recognized as legal. This different degree of accuracy reflects the more compelling task of assessing new rather then well known phenomena. The enforcer when investigating the effects of the new action receives a signal $\sigma = \{b, g\}$ on the state of nature, i.e. on the social consequences of the new action. With probability α_I the signal is incorrect when the true state of the world is the good one: in this case the enforcer considers action a as unlawful when the good state occurs, committing a "type-I error", or false negative.

articulated rule would consider resale price maintenance as illegal when adopted by a dominant firm, where this latter is identified by certain thresholds for market shares (x_1), entry conditions (x_2) and demand elasticity. (x_3), and when sales effort activities are irrelevant (x_4). On per-se rules v. effect-based rules se Gual et al. (2005) and Katsoulakos and Ulph (2009).

[8] This assumption is made only for clarity. We could consider both endogenous monitoring and accuracy. However, the main insight can be obtained also in the present setting with exogenous and perfect monitoring. We thank Patrick Rey for this suggestion.

Conversely, with probability α_{II} the signal is incorrect when the true state is the bad one, and a "type-II error" occurs, i.e. the enforcer will fail to identify a as unlawful when the true state is the bad one, an example of false positive. Hence,

$$\alpha_I = \Pr(\sigma = b\,|s = g) \quad \text{and} \quad \alpha_{II} = \Pr(\sigma = g\,|s = b).$$

We assume that the signals received are informative, i.e. $\alpha_I \leqslant \frac{1}{2}$ and $\alpha_{II} \leqslant \frac{1}{2}$.[9] Then type-$i$ accuracy, $i = I, II$, is measured by the probability of properly assessing the true state of the world $1 - \alpha_i$. Notice that since the signals are informative accuracies range in the interval $\left[\frac{1}{2}, 1\right]$, with the lowest level equal to $\frac{1}{2}$.

The level of accuracy can be improved by committing additional resources to obtain a more precise assessment of the effects. As we argued in the introduction, accuracy can be refined decreasing type-I, type-II or both types of errors. By adopting different protocols of investigation and standards of proof and by committing more resources, the legislator can reduce selectively type-I or type-II errors or he can symmetrically improve the assessment reducing both types of errors.[10]

We assume that the cost of obtaining a given type-i accuracy $i = I, II$, is increasing and convex in accuracy $1-\alpha_i$. More precisely, for type-i accuracy,$i = I, II$, the cost for the enforcer is $g_i(1 - \alpha_i)$, with $g_i' > 0$ and $g_i'' > 0$ for $\alpha_i \in \left[0, \frac{1}{2}\right]$, and with $g_i(\frac{1}{2}) = g_i'(\frac{1}{2}) = 0$ and $\lim_{\alpha_i \to 0} g_i(.) = \lim_{\alpha_i \to 0} g_i'(.) = \infty$.

In our setting, the firm has to pay a fine f with probability α_I if the new action is socially beneficial, and with probability $1-\alpha_{II}$ when it is socially harmful. Given the norm, we may have two different cases, according to the policy chosen by the enforcer. If she does not invest in accuracy, $\alpha_I = 1 - \alpha_{II} = \frac{1}{2}$ and the new action is treated in the same way in both states of the world. This is an instance of per-se legality (if combined with $f = 0$) or per se illegality (when $f > 0$) rule. When instead the enforcer improves accuracy we have $\alpha_I < 1 - \alpha_{II}$ and the fine is more likely when the action is socially harmful. The norm is then applied according to a discriminating (effect-based) legal standard.

Given the chosen level of accuracy, the enforcer evaluates the action according to a signal $\sigma = b, g$ (possibly wrong), and imposes a fine according to the perceived state of the world and the action observed. The fine will be zero following a status-quo action, i.e., $f(a_0) = 0$

[9] Assuming that, absent any accuracy effort, the probability of either error being $\frac{1}{2}$ is purely conventional. We might, for instance, assume that a minimum level of accuracy (a maximum probability of error $\overline{\alpha} < \frac{1}{2}$) must be ensured in order to avoid an overruling in the appeal phase. In this case all our arguments would apply to the additional resources devoted to improve accuracy beyond that minimum level, while the cost of accuracy would have a fixed component.

[10] As Katsoulakos and Ulph correctly observe, accuracy depends also on the legal standard adopted. If socially damaging actions are consided as illegal by the norm, and a per-se rule is adopted to enforce the norm, it is more likely to commit errors compared with an effect-based (discriminating) legal standard.

(with no error). If instead, the observed action is a and the signal suggest the bad state, the selected action a is considered unlawful and it is fined, i.e., $f(a\,|\sigma = b) = f$. Conversely, if the good state is signalled, a is lawful and no fine is set, i.e., $f(a\,|\sigma = g) = 0$.

The timing of the model is as follows. At time 0 nature chooses the state of the world $s = \{g, b\}$ which is not observed by any agent. Agents know that the probability of the bad state is $\beta > 0$. At time 1, given the legal rule, the enforcer commits to a certain fine f and to accuracies α_I and α_{II}. At time 2, having observed the policy set by the enforcer, the firm chooses the research investment I and learns with probability $p(I) = I$ how to implement the new action a and the state of the world. At time 3, the firm chooses an action, conditional on what it learnt in the previous stage. Finally, at time 4 the action chosen determines the private profits and the social welfare; the enforcer collects evidence (with errors) and possibly levies a fine.

3 First best

As the opportunities created by research generate either positive or negative externalities, depending on the state of nature, public policy may be beneficial. To evaluate the benefits of public interventions, it is useful to benchmark them against the first-best outcome (FB), which would obtain if the enforcer could control firms' choices I and a directly.

Let us denote by a_s with $s = b, g$ the action chosen by the firm — if research is successful — in the bad and in the good state, respectively. If the research activity is successful, unconstrained welfare maximization calls for the new action in the good state and the status-quo action in the bad state, that is $a_g = a$ and $a_b = a_0$. For given level of the innovative investment I, the first-best expected welfare is therefore

$$\mathrm{E}(W_{FB}) = I(1-\beta)\overline{W} - c\frac{I^2}{2}. \tag{2}$$

The first-order condition with respect to I yields the corresponding investment level

$$I_{FB} = \frac{(1-\beta)\overline{W}}{c}, \tag{3}$$

which is increasing in the likelihood of the good state $(1-\beta)$ and in the associated welfare gain $\overline{W}$, and decreasing in the marginal cost of innovative activity c. Then, the expected welfare evaluated at the first-best policies is

$$\mathrm{E}(W^*_{FB}) = \frac{(1-\beta)^2\overline{W}^2}{2c}.$$

In what follows, the policy maker is assumed not to control firms' choices directly, but to influence them via penalties: firms are free to implement their preferred action, but they are aware that public intervention may occur *ex post* in the form of fines, whenever social harm is assessed to have occurred (with errors).

4 Firm's choices: actions and research

At stage 3, depending on whether its research was successful or not, the firm chooses an action. Two elements affect the selection of the profit maximizing action, namely the outcome of the research investment and the enforcement policy. If the innovative activity has been unsuccessful, the only feasible choice is the status-quo action a_0 with associated profits $\Pi(a_0) = 0$ and welfare $W(a_0) = 0$. If instead the innovative activity has been successful, the firm is able to take also the new action a, and in this case the enforcement policy affects the firm's choice.

In the good state ($s = g$) the action a is not socially harmful and therefore, if correctly assessed, it is considered lawful. Nevertheless, with probability α_I the authority's perceived state of the world is the bad one ($\sigma = b$) and the action is erroneously judged unlawful. Then, when $a_g = a$ the expected profits are equal to $\Pi - \alpha_I f$. If, alternatively, the firm chooses the action $a_g = a_0$, profits are equal to 0 and there is no error in enforcement.

Turning to the bad state ($s = b$), the new action a, if correctly recognized, is socially harmful, and therefore unlawful, but the fine is levied only with probability $(1 - \alpha_{II})$ since with probability α_{II} the enforcer receives the wrong signal $\sigma = g$. In this case when the firm chooses $a_b = a$, the expected profits are equal to $\Pi - (1 - \alpha_{II}) f$. To break the ties, we assume that when indifferent the firm chooses the welfare maximizing action. Hence, when $\Pi - \alpha_I f = 0$, since the state is good, the firm chooses a. Instead, if $\Pi - (1 - \alpha_{II}) f = 0$, being in the bad state, the firm will select the status-quo action a_0. Finally, it is simple to check that $\frac{\Pi}{\alpha_I} \geqslant \frac{\Pi}{(1-\alpha_{II})}$ or, equivalently, $\alpha_I + \alpha_{II} \leqslant 1$.[11]

It is interesting to notice that more type-I accuracy, i.e. a smaller α_I, makes the choice of the new action a more convenient in the good state, while more type-II accuracy, i.e. a smaller α_{II}, makes the choice of the new action less attractive in the bad state. This observation is consistent with the received view that accuracy exerts a positive impact on the ability of the legislator to implement the desired firm's actions, what is usually called marginal deterrence.

The next Lemma directly follows from the discussion above, and shows that according to the outcome of research (successful or unsuccessful), and according to the policy chosen by the legislator (the level of the fine f and the type-I and type-II errors α_I and α_{II}), only four action choices can be implemented:

[11] In fact, $\frac{\Pi}{(1-\alpha_{II})} > \frac{\Pi}{\alpha_I}$ implies $\alpha_I + \alpha_{II} > 1$, which is impossible for $\alpha_i \in \left[0, \frac{1}{2}\right]$.

Lemma 1(Actions Choice): *If research is successful three cases may arise:*

(i) (status-quo action - SQ): for $f > \frac{\Pi}{\alpha_I}$*, the firm chooses the status quo action in both states, i.e.* $a_g = a_b = a_0$;

(ii) (efficient actions - EA): for $\frac{\Pi}{\alpha_I} \geqslant f \geqslant \frac{\Pi}{(1-\alpha_{II})}$*, the firm chooses the status quo action in the bad state and the new action in the good state, i.e.* $a_g = a$ and $a_b = a_0$;

(iii)(new action - NA): for $f < \frac{\Pi}{(1-\alpha_{II})}$*, the firm chooses the new action in both states, that is* $a_g = a_b = a$;

(iv) If reearch is unsuccessful, the firm chooses the status quo action a_0.

At stage 2, knowing the accuracies and the fine, the firm chooses its research investment I so as to maximize its expected profits, anticipating the optimal actions that it will choose at stage 3. The firm learns how to carry out the new project with probability $p(I) = I$ and its expected profits at this stage are:

$$E\Pi = I\left\{\beta\left[\max\left\{\Pi - (1-\alpha_{II})\, f, 0\right\}\right] + (1-\beta)\left[\max\left\{\Pi - \alpha_I f, 0\right\}\right]\right\} - c\frac{I^2}{2}, \qquad (4)$$

where the first term is the expected gain from research (net of the expected fines), and the second term is the cost of research. Notice that, since the investment is chosen before observing the true state of nature, the expected profits depend on both the policies adopted when the new action is lawful and when it is instead unlawful. The profit maximizing research investment is described in the following result for the three different action choices described in Lemma 1.

Lemma 2 (Investment choice): *The research investment I chosen by the firm in the three different cases is*

$$\widehat{I}(f, \alpha_I, \alpha_{II}) = \begin{cases} (SQ) & 0 < I_{FB} & \text{for } f > \frac{\Pi}{\alpha_I}, \\ (EA) & \frac{(1-\beta)(\Pi - \alpha_I f)}{c} < I_{FB} & \text{for } \frac{\Pi}{\alpha_I} \geqslant f \geqslant \frac{\Pi}{(1-\alpha_{II})}, \\ (NA) & \frac{\Pi - [\beta(1-\alpha_{II}) + (1-\beta)\alpha_I] f}{c} \lesseqqgtr I_{FB} & \text{for } f < \frac{\Pi}{(1-\alpha_{II})}. \end{cases} \qquad (5)$$

In the first two cases the investment is lower than the first best while in case (iii) we may have under or over-investment.

Notice that due to assumption (1) $\widehat{I}(.) \in [0, 1]$. Moreover, in the efficient action case (EA) only type-I errors α_I affect the investment, since in the bad state the firm chooses the status-quo action which is correctly assessed by the legislator. Less accuracy (a higher α_I), by increasing over-deterrence, reduces the expected profits and the investment in innovative activity. In the new action case (NA), instead, the firm always chooses the new action a and the investment is discouraged by more over-deterrence (higher α_I) or less under-deterrence (lower α_{II}). In this case, therefore, the two accuracies affect the innovative investment in opposite directions.

By direct comparison of (5) with (3) it is immediate to see that there is always underinvestment in research in the status-quo and efficient action cases, while in the new action case it is possible to have both under or overinvestment. Then, case EA — with the enforcer that implements the efficient action in each state — is not necessarily the preferred outcome once we take into account the insufficient level of research investment. Indeed, the underinvestment result, together with the cost of accuracy, explains why the efficient action case will not be always the optimal one.

5 Optimal enforcement policies

After having identified the firm's choice of action and research, we study the optimal policy in two steps: first, we identify the optimal fine and accuracies for any implemented course of actions, namely the status quo, the efficient and the new action policies. Then, we compare welfare in the different scenarios selecting the overall optimal policy.

In the present setting, the legislator controls three instruments: α_I and α_{II}, which are costly and the fine f, that is a pure transfer. The expected welfare, taking into account the firm's optimal choices, is:

$$EW = \begin{cases} (SQ) \quad 0 - g_I(.) - g_{II}(.) - 0 & \text{for } f > \frac{\Pi}{\alpha_I}, \\ (EA) \quad \widehat{I}(.)(1-\beta)\overline{W} - g_I(.) - g_{II}(.) - c\frac{\widehat{I}(.)^2}{2}, & \text{for } \frac{\Pi}{\alpha_I} \geqslant f \geqslant \frac{\Pi}{(1-\alpha_{II})}, \\ (NA) \quad \widehat{I}(.)\Delta W(\beta) - g_I(.) - g_{II}(.) - c\frac{\widehat{I}(.)^2}{2}, & \text{for } f < \frac{\Pi}{(1-\alpha_{II})}, \end{cases} \tag{6}$$

where $\Delta W(\beta) \equiv [\beta\underline{W} + (1-\beta)\overline{W}]$ is the expected welfare change due to the new action a, while the last three terms correspond to the public costs of accuracies and the private costs of the research investment.

5.1 The status quo policy (SQ)

The first regime, identified by the condition $f > \frac{\Pi}{\alpha_I}$, is also the simplest to study. In this case the legislator sets the fine at such a high level that, in case of successful innovative activity, the firm prefers to stick to the status-quo action a_0 in both states. Indeed, due to the errors the fine is given in both states, and the expected profits from the innovative investment are non positive, inducing the firm not to invest in research, i.e. $\widehat{I} = 0$. The cheapest way to implement this regime is by choosing no accuracy and a sufficiently high fine, that is, $\alpha_I^* = \alpha_{II}^* = \frac{1}{2}$ and $f > 2\Pi$. The expected welfare is then

$$\text{EW}^*_{SQ}(\beta) = 0, \text{ for any } \beta.$$

The status-quo outcome is consistent with a legal standard that always prohibits the new action and fines it with the same probability (equal to $\frac{1}{2}$) no matter what are its social consequences, what we could call a per-se illegality rule.

5.2 The efficient action policy (EA)

In the efficient action regime fine and accuracies are set so as to implement the welfare maximizing actions, that is, $a_g = a$ in the good state and $a_b = a_0$ in the bad state. However, as mentioned in Lemma 2, this positive effect is counterbalanced by a level of investment in research that is too low compared to the first best for any given probability of the bad state β. In other words, in this regime the enforcer's optimal policy aims at reaching both marginal deterrence and the right incentives to innovate. The enforcer uses the policy instruments to solve the following problem:

$$\max \widehat{I}(.)(1-\beta)\overline{W} - g_I(1-\alpha_I) - g_{II}(1-\alpha_{II}) - c\frac{\widehat{I}(.)^2}{2}, \tag{7}$$

$$\text{s.t.} \quad : \quad \frac{\Pi}{\alpha_I} \geqslant f \qquad \text{and} \qquad f \geqslant \frac{\Pi}{(1-\alpha_{II})}.$$

Recalling that in this regime there is always underinvestment in innovative activity, it is immediate to see that the second constraint must be binding. Since from Lemma 2, $\widehat{I} = \frac{(1-\beta)(\Pi-\alpha_I f)}{c}$, lowering f makes the first constraint slacker while still inducing the welfare maximizing actions, and mitigates the underinvestment in research more cheaply than through an increase in type-I accuracy α_I. Then, $f^*(a\,|\sigma = b) = \frac{\Pi}{(1-\alpha_{II})}$ and

$$\widehat{I} = \frac{\Pi(1-\beta)}{c}\left(1 - \frac{\alpha_I}{1-\alpha_{II}}\right).$$

We assume that the cost of accuracy are sufficiently convex. Then, the following Proposition, proved in the Appendix, establishes the optimal policy in the efficient action regime.

Proposition 1 (The efficient action policy): *The optimal policy entails an effect-based legal standard with* $\alpha_I^* \in \left(0, \frac{1}{2}\right)$, $\alpha_{II}^* \in \left(0, \frac{1}{2}\right)$ *and* $f = \frac{\Pi}{(1-\alpha_{II}^*)} \in (\Pi, 2\Pi)$ *for* $\beta \in (0,1)$ *and a per-se illegality legal standard with* $\alpha_I^* = \alpha_{II}^* = \frac{1}{2}$ and $f = 2\Pi$ for $\beta = 1$. *Moreover, if* $g_I(.) = g_{II}(.) = g(.)$ *then* $\alpha_I^* < \alpha_{II}^*$ for $\beta \in (0,1)$. *The expected welfare is equal to*

$$EW_{EA}(\beta) = \frac{\Pi\,(1-\beta)^2}{c}\left(1 - \frac{\alpha_I^*}{1-\alpha_{II}^*}\right)\left(\overline{W} - \frac{\Pi}{2}(1 - \frac{\alpha_I^*}{1-\alpha_{II}^*})\right) - g_I(.) - g_{II}(.)$$

and is decreasing and concave in the probability of the bad state β.

It is interesting to notice that all the instruments — both types of accuracy as well as the fine f — are used. To understand the different effects at work it is useful to discuss the result in detail. The level of investment in the efficient action case is determined by the fine f and the type-I error α_I. Since the investment is below the first-best, the legislator is willing to boost it by reducing the fine f at the lowest level consistent with the implementation of the efficient action a_0 in the bad state, i.e., at $f = \frac{\Pi}{(1-\alpha^*_{II})}$. This corner solution suggests that there is a tension, in the optimal policy design, between implementing the efficient actions, a goal that calls for a sufficiently high fine, and giving the proper incentives to the innovative activity, an objective that instead would require to reduce the fine. For this reason, reducing the fine at $f = \frac{\Pi}{(1-\alpha^*_{II})}$ is not enough, and the legislator improves both type of accuracies to further boost the investment while still implementing the efficient actions.

Hence, both accuracies are implemented, acting as complements in the optimal policy. Notice that since the enforcer invests in accuracy the fine is given with different probabilities α^*_I and $1-\alpha^*_{II}$ in the two states of the world, implying that the legal standard discriminates between socially beneficial and socially damaging actions, an instance of effect-based legal standard.

Reducing type-I errors α_I, has a direct positive effect on the investment that is stronger than the second, indirect, one through α_{II}, explaining the asymmetric level of investment in accuracies. This latter result calls for an asymmetric protocol of investigation, that allows to reduce differently type I and type II errors. To illustrate this point, consider the following example, drawn from antitrust policy. Suppose that the welfare effects of a given practice of a dominant firm depend, in decreasing order, on four fundamental variables: market shares, entry conditions, demand elasticity and cost efficiencies. If we just consider the action chosen (a_0 or a) with no additional element, we opt for no accuracy ($\alpha_I = \alpha_{II} = \frac{1}{2}$) and just conclude for illegality if we observe the new action, an instance of (form-based) per-se illegality.

If we instead opt for greater accuracy, we might proceed in different ways in assessing additional variables, improving accuracy selectively on one or the other type of error. A case of selective accuracy can be illustrated by the following protocol of investigation. 1) The enforcer considers first the market share of the dominant firm: if this is below a certain threshold the enforcer clears the case, if instead the market share is above the threshold the investigation proceeds to the second step (entry conditions); 2) if entry is easy the case is cleared, while difficult entry leads to the third step (demand elasticity); 3) If the demand is elastic the case is cleared while inelastic demand required to move to the last step (efficiencies); 4) substantial efficiencies call for clearing the case while low cost reductions determine a negative judgement.

This example illustrates a protocol of investigation in which further levels of accuracy are implemented asymmetrically, requiring a more compelling standard of proof every time a negative signal $\sigma = b$ is assessed, while the firm is judged innocent if a positive signals is obtained at any stage of the investigation. This way the probability of condemning an innocent firm (type-I error) is reduced much more than that of acquitting a guilty one.[12].

5.3 The new action policy (NA)

In the new action regime the legislator chooses a fine low enough so that the profits from the new action exceed the maximum fine, implying that, if the innovative effort is successful, the firm always prefers to choose the new action. Although in this case the legislator renounces to affect the firm's choice of actions, a role for deterrence remains through the policy effect on the research investment

$$I = \frac{\Pi - [\beta\,(1-\alpha_{II}) + (1-\beta)\alpha_I]f}{c}.$$

The relevant first order conditions are:

$$\frac{\partial EW}{\partial f} = [\Delta W(\beta) - c\widehat{I}]\frac{\partial \widehat{I}}{\partial f} \gtreqless 0, \tag{8}$$

$$\frac{\partial EW}{\partial \alpha_I} = [\Delta W(\beta) - c\widehat{I}]\frac{\partial \widehat{I}}{\partial \alpha_I} + g_I' \geqslant 0, \tag{9}$$

and

$$\frac{\partial EW}{\partial \alpha_{II}} = [\Delta W(\beta) - c\widehat{I}]\frac{\partial \widehat{I}}{\partial \alpha_{II}} + g_{II}' \geqslant 0, \tag{10}$$

where $\frac{\partial \widehat{I}}{\partial f} = -\frac{\beta(1-\alpha_{II})+(1-\beta)\alpha_I}{c} < 0$, $\frac{\partial \widehat{I}}{\partial \alpha_I} = -\frac{(1-\beta)f}{c} \leqslant 0$ and $\frac{\partial \widehat{I}}{\partial \alpha_{II}} = \frac{\beta f}{c} \geqslant 0$. The three derivatives have the same structure. The first term captures the marginal effect of the policy variables on expected welfare, through the impact on the research investment. The second term, which is zero in the case of fines, is the marginal cost of the policy. The optimal choice of the policy variables, therefore, depends on the sign of the marginal social value of the research investment, $\Delta W(\beta) - c\widehat{I}$ that can be positive or negative. Hence, in general we may expect internal as well as corner solutions.

To characterize the optimal policy, we substitute the optimal investment $\widehat{I}$ chosen by the firm according to (5) in $[\Delta W(\beta) - c\widehat{I}] = 0$, and we solve for β. The result is the following set of loci:

$$\widehat{\beta}(f, \alpha_I, \alpha_{II}) = \frac{\overline{W} - \Pi + \alpha_I f}{\overline{\overline{W}} - \underline{W} + \alpha_I f - (1-\alpha_{II})\,f},$$

which describe, for given policy parameters $(f, \alpha_I, \alpha_{II})$, the probability of the bad state β that makes the marginal social value of the research investment equal to zero. The

[12] Different levels of type-I accuracy may be obtained then by interrupting the protocol of investigation at stage 3 or 2.

equilibrium analysis focusses on two relevant loci corresponding to different combinations of policy parameters:

$$\beta_0 = \widehat{\beta}(0, \frac{1}{2}, \frac{1}{2}) = \frac{\overline{W} - \Pi}{\overline{W} - \underline{W}},$$

along which the research investment is welfare neutral when $f = 0$ and $\alpha_I = \alpha_{II} = \frac{1}{2}$, and

$$\beta_1 = \widehat{\beta}(2\Pi, \frac{1}{2}, \frac{1}{2}) = \frac{\overline{W}}{\overline{W} - \underline{W}},$$

corresponding to the policy parameters $f = 2\Pi$ and $\alpha_I = \alpha_{II} = \frac{1}{2}$. This latter policy combination induces the firm not to invest in research, i.e. $\widehat{I} = 0$. Moreover, at $\beta = \beta_1$ the new action a gives a zero welfare gain, that is $\Delta W(\beta_1) = 0$. Finally, notice that $1 > \beta_1 > \beta_0 > 0$.

The following Proposition, proved in the Appendix, characterizes the optimal policy which implements the NA regime for different likelihood of the bad state β.

Proposition 2 (The new action policy): *We can distinguish the following regions:*

- *If $\beta \in [0, \beta_0]$ the optimal policy entails a per-se legality rule (laissez-faire) with $f = 0$ and $\alpha_I^* = \alpha_{II}^* = \frac{1}{2}$. The expected welfare is strictly positive and equal to $EW_{NA}(\beta) = \frac{\Pi}{c}\left(\Delta W(\beta) - \frac{\Pi}{2}\right) > 0$;*
- *If $\beta \in (\beta_0, \beta_1)$ the optimal policy adopts a per-se illegality rule and requires to set $\alpha_I^* = \alpha_{II}^* = \frac{1}{2}$ and*

$$f^* = 2\left[\Pi - \Delta W(\beta)\right] \in (0, 2\Pi).$$

The expected welfare is strictly positive and equal to $EW_{NA}(\beta) = \frac{\Delta W(\beta)^2}{2c} > 0$; for $\beta \to \beta_1$ we have $\widehat{I}(\beta) = \frac{\Delta W(\beta)}{c} \to 0$ and $EW_{NA}(\beta) = \frac{\Delta W(\beta)^2}{2c} \to 0$.

In the new action regime the enforcer renouces to elicit the efficient action and concentrates his efforts on affecting the research investment. She does not invest in accuracy and fines with the same probability 1/2 both beneficial and harmful action. Hence, the enforcer in this case adopts per-se legal standards. When the likelihood of the bad state β is small, i.e. below β_0, the research investment is socially desirable. The optimal policy requires that the new action is not penalized ($f = 0$) even if ex-post it is found to be welfare decreasing in the bad state. The *ex-ante* positive effect on research incentives, indeed, more than counterbalances the *ex-post* limited losses. Accuracy in this case is irrelevant because the firm does not pay any fine no matter what is the final decision. This policy is equivalent to *laissez-faire* or a per-se legality legal standard.

Once we move to a region where social losses are more likely, between β_0 and β_1 a *laissez-faire* regime is too lax and we want to marginally limit the research investment. This result is obtained using the cheapest tool, i.e. increasing the fine without spending on accuracy. The new action is therefore sanctioned with the same probability $(\frac{1}{2})$ in both states of the world, a case of per-se illegality legal standard. In this case the new action leads to a positive but decreasing expected welfare variation, i.e. $\Delta W(\beta) > 0$, while the marginal social value of the research investment, taking into account its costs, is kept at zero by progressively reducing I. At the upper bound β_1, the innovative investment is completely discouraged, and the firm never learns how to implement the new action. In this case, the welfare remains at the status-quo (zero) level and the issue of marginal deterrence, that is the ability to influence the ex-post choice of the action, becomes irrelevant.

5.4 Optimal policy with respect to β

The three cases considered involve a different use of the instruments available and in the (implicit) legal standard adopted. If the legislator favors the goal of implementing the efficient actions (EA), a combination of fines and accuracies is needed, with a bias in favor of type-I accuracy. In this case, indeed, the policy tries to affect at the same time the choice of the innovative activity and of the actions, so that more tools are required. Moreover, the norm is applied differently in the good and the bad state, implying a discrimination consistent with an effect-based legal standard. When instead the enforcer renounces to affect the choice of the actions (NA) and focusses only on the choice of the investment, a single tool suffices, and she uses the costless fines without promoting accuracy. Finally, if the new action must be avoided (SQ) fines are set at a sufficiently high level with the aim of inducing the firm to refrain from investing in innovative activity and no resources are wasted to increase accuracy. In the latter two cases, the policy acts in the same way whether the action is socially beneficial or harmful, and is therefore equivalent to a per-se (legality or illegality) legal standard.

Once identified the policies consistent with the implementation of different firm's choices, we want to compare their expected welfare levels and find out the overall optimal policy according to different levels of β, the probability that the new action is socially harmful. The following Proposition, proved in the Appendix, establishes the optimal policy for different values of β:

Proposition 3 (Overall optimal policy): *$\exists$ a $\widetilde{\beta} < \beta_1$ such that the NA policy is preferred to the EA one iff $\beta \leqslant \widetilde{\beta}$ and vice-versa. The SQ policy is never chosen.*

The previous Proposition shows that when the new action is unlike to give social losses, the enforcer focusses on the research incentives promoting *laissez-faire* or moderate fines, without investing in accuracy. Conversely, for a higher likelihood of social harm, the enforcer prefers to elicit the efficient actions focussing on marginal deterrence even at the cost of depressing firm's research investment. This outcome is obtained by a combination of fines and asymmetric accuracies, with a stronger effort to reduce type-I error than type-II errors. When the probability of harmful innovation is very high, the policy reduces the research investment to zero, implementing in the limit the status-quo.

Our result has implications also in terms of the legal standard that is adopted by the enforcer, an issue analyzed in Katsoulakos and Ulph (2009). For $\beta \in [0, \beta_0]$ the new action is not sanctioned in either state of the world, and is therefore treated equivalently to a per-se legality rule. For $\beta \in (\beta_0, \widetilde{\beta}]$ the new action is sanctioned with increasing fines and with the same probability (1/2) in both states of the world, without discriminating its social effects. Hence, in this region the policy is equivalent to a per-se illegality rule enforced with increasing fines. Once we enter into the efficient action region, for $\beta \in (\widetilde{\beta}, 1)$ accuracy implies a different probability of fines in the two states of the world. This policy is therefore consistent with a discriminating (effect-based) rule. Finally, for $\beta = 1$ no accuracy is adopted and we are back to a per-se illegality rule.

6 Conclusions

In the traditional model of law enforcement a private agent selects an illegal action from a set of privately convenient but socially damaging actions by comparing their expected benefits and fines. Marginal deterrence, in this setting, is the key effect.

In this paper, the agents first have to invest resources in learning – what we call research – and then, if successful, they are able to choose a new action that *ex post* may be welfare enhancing or reducing (illegal). The optimal policy, determining the probability of being fined *ex post*, affects both the expected profits from the new action and the *ex-ante* incentives to invest in research. In our setting the instruments of the legislator are the level of fines and the level of accuracy in assessing the social consequences (legality) of the actions.

The policy affects the firm's choices on two grounds: by influencing the ex-post choice of the action, an instance of marginal deterrence, and through the ex-ante incentives to research investment. The two goals are not always aligned, and the relative importance of the two effects shapes the optimal policy.

We have shown that the likelihood of social harm β drives the design of the optimal policy: for a sufficiently low probability β, the enforcer does not invest in accuracy and focusses on affecting the research incentives through zero (per-se legality legal standard) or low fines (moderately enforced per-se illegality rule), allowing the firm to choose the new action in any state and disregarding marginal deterrence. Conversely, for a higher likelihood of social harm β, the enforcer prefers to improve marginal deterrence by eliciting the ex-post efficient actions at the cost of depressing firm's research investment. This is obtained through a combination of fines and accuracies with a bias in favor of type-I accuracy and the adoption of asymmetric protocols of investigation. By investing in accuracy the enforcer sanctions the new action with different probabilities when harmful or beneficial, adopting a discriminating (effect-based) legal standard. In the limit, when the social harm of innovation is almost certain, the policy completely displaces the research investment through a very high fine in any contingency (per-se illegality legal standard).

Appendix

Proof of Proposition 1. Substituting, $f = \frac{\Pi}{(1-\alpha_{II})}$ in (7) and maximizing with respect to α_I and α_{II} we get the following first order and complementary slackness conditions, where λ is a Lagrange multiplier:

$$g_I' = A + \frac{\lambda\Pi}{\alpha_I^2} \geqslant 0, \tag{11}$$

$$g_{II}' = \frac{\alpha_I}{(1-\alpha_{II})}\left[A + \frac{\lambda\Pi}{(1-\alpha_{II})^2}\right] \geqslant 0, \tag{12}$$

$$\lambda\left(\frac{\Pi}{\alpha_I} - \frac{\Pi}{(1-\alpha_{II})}\right) = 0, \tag{13}$$

where

$$A \equiv \frac{\Pi(1-\beta)^2}{c(1-\alpha_{II})}\left[\overline{W} - \Pi\left(1 - \frac{\alpha_I}{(1-\alpha_{II})}\right)\right] > 0$$

for $\beta < 1$ and $A = 0$ for $\beta = 1$.

Consider first the case where $\beta < 1$: the RHS in (11) and (12) are positive for any $\lambda \geqslant 0$, then the Inada conditions $g_i'(1/2) = 0$, and $\lim_{\alpha_i \to 0} g'(.) = \infty$ imply that $\alpha_I^* \in \left(0, \frac{1}{2}\right)$ and $\alpha_{II}^* \in \left(0, \frac{1}{2}\right)$. Moreover (13) implies that $\lambda = 0$. The second order conditions for α_I and α_{II} are:

$$\begin{aligned}
\frac{\partial^2 EW}{\partial \alpha_I^2} &= -\frac{\Pi^2(1-\beta)^2}{c(1-\alpha_{II})^2} - g_I'' < 0,\\
\frac{\partial^2 EW}{\partial \alpha_{II}^2} &= \frac{\Pi(1-\beta)^2\alpha_I}{c(1-\alpha_{II})^3}\left[\Pi\left(2 - \frac{3\alpha_I}{(1-\alpha_{II})}\right) - 2\overline{W}\right] - g_{II}'' < 0,
\end{aligned}$$

and the determinant of the Hessian matrix:

$$|H_2| = g_I'' g_{II}'' + \frac{\Pi(1-\beta)^2}{c(1-\alpha_{II})^2}\left\{\begin{array}{l} \left[2\overline{W} - \Pi\left(2 - \frac{3\alpha_I}{(1-\alpha_{II})}\right)\right]\left[g_I''\frac{\alpha_I}{(1-\alpha_{II})^2} + \frac{\Pi^2(1-\beta)^2\alpha_I}{c(1-\alpha_{II})^3}\right] + \\ +\Pi g_{II}'' - \frac{\Pi(1-\beta)^2}{c(1-\alpha_{II})^2}\left[2\overline{W} - \Pi\left(2 - \frac{3\alpha_I}{(1-\alpha_{II})}\right)\right]^2 \end{array}\right\} > 0,$$

that is positive when, g_I'' and g_{II}'' are sufficiently large (only the last term in the curly brackets is negative). Turning to the comparison of the levels of accuracy, when $A > 0$, i.e., for $\beta < 1$, since $\alpha_I^* + \alpha_{II}^* < 1$ and $\frac{\alpha_I^*}{(1-\alpha_{II}^*)} < 1$, the first order conditions imply that:

$$g_{II}'(1-\alpha_{II}^*) = \frac{\alpha_I^*}{\left(1-\alpha_{II}^*\right)} g_I'(1-\alpha_I^*) < g_I'(1-\alpha_I^*)\,(>0).$$

We then conclude that when the two types of accuracy entail the same costs, i.e. $g_I(.) = g_{II}(.) = g(.)$ we have $1 - \alpha_{II}^* < 1 - \alpha_I^*$ and therefore $\alpha_I^* < \alpha_{II}^*$ for $\beta < 1$.

Consider now the case $\beta = 1$ and $A = 0$. Then, λ must be zero as well. We prove this by contradiction. Suppose that $\lambda > 0$. In this case (11) and (12) are positive, implying that $\alpha_I^* + \alpha_{II}^* < 1$ so that the first constraint does not bind. Then, in the complementary slackness condition the term in parenthesis is positive and $\lambda = 0$. Hence, for $\beta = 1$ we have $A = \lambda = 0$ and $\alpha_I^* = \alpha_{II}^* = \frac{1}{2}$.

Finally, differentiating the expected welfare

$$EW_{EA}(\beta, \alpha_I^*(\beta), \alpha_{II}^*(\beta)) = \frac{\Pi(1-\beta)^2}{c}\left(1 - \frac{\alpha_I^*}{1-\alpha_{II}^*}\right)\left(\overline{W} - \frac{\Pi}{2}(1 - \frac{\alpha_I^*}{1-\alpha_{II}^*})\right) - g_I(.) - g_{II}(.)$$

with respect to β,we get:

$$\frac{dEW}{d\beta} = \frac{\partial EW}{\partial \beta} + \frac{\partial EW}{\partial \alpha_I}\frac{\partial \alpha_I^*}{\partial \beta} + \frac{\partial EW}{\partial \alpha_{II}}\frac{\partial \alpha_{II}^*}{\partial \beta},$$

where the first term (direct effect) is negative and the last two terms are zero due to the FOC (envelope theorem). Differentiating one more time and taking into account that $\frac{\partial EW}{\partial \alpha_I} = \frac{\partial EW}{\partial \alpha_{II}} = 0$ we get:

$$\frac{d^2EW}{d\beta^2} = \frac{\partial^2 EW}{\partial \beta^2} + \frac{\partial^2 EW}{\partial \alpha_I^2}\left(\frac{\partial \alpha_I^*}{\partial \beta}\right)^2 + \frac{\partial^2 EW}{\partial \alpha_{II}^2}\left(\frac{\partial \alpha_{II}^*}{\partial \beta}\right)^2 + 2\frac{\partial^2 EW}{\partial \alpha_I \partial \alpha_{II}}\left(\frac{\partial \alpha_I^*}{\partial \beta}\frac{\partial \alpha_{II}^*}{\partial \beta}\right),$$

where $\frac{\partial^2 EW}{\partial \beta^2} > 0$, $\frac{\partial^2 EW}{\partial \alpha_I^2} < 0$, $\frac{\partial^2 EW}{\partial \alpha_{II}^2} < 0$ and $\frac{\partial^2 EW}{\partial \alpha_I \partial \alpha_{II}} < 0$. Moreover $\frac{\partial \alpha_I^*}{\partial \beta} > 0$ and $\frac{\partial \alpha_{II}^*}{\partial \beta} > 0$ if g_i'' are sufficiently large. Hence, the first term is positive but the other three terms are negative. The expected welfare is then always decreasing in β and concave when g_i'' are sufficiently positive and large so to neutralize the only positive term. ■

Proof of Proposition 2. Consider the first order conditions (8), (9) and (10).

If $\beta \in [0,\ \beta_0)$, when $f = 0$ and $\alpha_I = \alpha_{II} = \frac{1}{2}$ the marginal social value of the innovative activity $[\Delta W(\beta) - c\widehat{I}]$ is positive and therefore $\frac{\partial EW}{\partial f} < 0$: in this region increasing the fine f reduces the incentive to innovate and the expected welfare. Consequently, the optimal policy is to set the fine $f^* = 0$, i.e. not to penalize the new action. In this case $\frac{\partial \widehat{I}}{\partial \alpha_I} = \frac{\partial \widehat{I}}{\partial \alpha_{II}} = 0$ and therefore $\frac{\partial EW}{\partial \alpha_I} = g_I'$ and $\frac{\partial EW}{\partial \alpha_{II}} = g_{II}'$. Consequently, it is optimal to set $\alpha_I^* = \alpha_{II}^* = 1/2$ making $g_I'(\frac{1}{2}) = g_{II}'(\frac{1}{2}) = 0$. This outcome is therefore equivalent to the one arising if the new action was lawful, i.e. a *laissez-faire* or a *per-se legality* rule. Notice that the optimal policy parameters are indeed those corresponding to the threshold β_0: this ensures that below β_0, the marginal social value of innovative activity is indeed positive when the optimal policy is chosen. At $\beta = \beta_0$ the same argument holds with $[\Delta W(\beta) - c\widehat{I}] = 0$. Substituting the optimal policy in (6) we get $EW_{NA}(\beta) = \frac{\Pi}{c}\left(\Delta W(\beta) - \frac{\Pi}{2}\right)$ that is decreasing and linear in β. At the upper bound $\beta = \beta_0$ we have $\widehat{I} = \frac{\Pi}{c}$ and $\Delta W(\beta) = c\widehat{I} = \Pi$. Substituting in the expected welfare for $\beta = \beta_0$ we get $EW_{NA}(\beta_0) = \frac{\Pi^2}{c} > 0$. Hence, $EW_{NA}(\beta) > 0$ for any $\beta \in [0,\ \beta_0]$.

If $\beta \in (\beta_0,\ \beta_1)$, once we move above β_0 for $f = 0$ the marginal social value of innovative activity becomes negative and $\frac{\partial EW}{\partial f} > 0$, suggesting to increase the fine. Notice that a higher fine affects $[\Delta W(\beta) - c\widehat{I}]$ reducing the investment $\widehat{I}$ and its marginal cost and increasing the marginal social value of innovation. The optimal fine f^* is determined by the condition $\Delta W(\beta) - c\widehat{I}(f^*) = 0$. When $f = f^*$ the first order conditions for α_I^* and α_{II}^* are solved for $\alpha_I^* = \alpha_{II}^* = 1/2$ since $g_I'(\frac{1}{2}) = g_{II}'(\frac{1}{2}) = 0$. Substituting the expressions of $\Delta W(\beta)$ and $\widehat{I}$ and solving we obtain

$$f^* = 2\left[\Pi - \beta \underline{W} - (1-\beta)\overline{W}\right].$$

The second order condition then is $\partial^2 EW/\partial f^2 = -1/4c < 0$, since the other variables are set at the corner solutions. Substituting the optimal policy in (5) and (6) we get $\widehat{I}(\beta) = \frac{\Delta W(\beta)}{c} > 0$ and $EW_{NA}(\beta) = \frac{\Delta W(\beta)^2}{2c} > 0$, which is greater than zero and decreasing for $\beta \in (\beta_0,\ \beta_1)$. Notice also that for $\beta \to \beta_1$ the innovative investment $\widehat{I}(\beta)$ and the expected welfare $EW_{NA}(\beta)$ tend to zero. In this case, the firm never learns how to implement the new action.

Finally, if $\beta \in [\beta_1, 1]$ the optimal policy is not determined. Notice in fact that f^* is increasing in β and tends to 2Π when $\beta \to \beta_1$. However, for such a high fine the firm would never choose the new action a if the innovative investment is successful. But with such a high fine no investment is exerted. ■

Proof of Proposition 3. Consider the expression of the expected welfare under the three policies:

$$\begin{aligned}
EW_{SQ}(\beta) &= 0 \text{ for } \forall\beta,\\
EW_{EA}(\beta) &= \frac{\Pi}{c}(1-\beta)^2\left(1-\frac{\alpha_I^*}{1-\alpha_{II}^*}\right)\left(\overline{W}-\frac{\Pi}{2}(1-\frac{\alpha_I^*}{1-\alpha_{II}^*})\right) - g_I(.) - g_{II}(.) \quad \text{for } \forall\beta,\\
EW_{NA}(\beta) &= \begin{cases} \frac{\Pi}{c}\left(\Delta W(\beta) - \frac{\Pi}{2}\right) & \text{for } \beta \in [0,\beta_0] \\ \frac{\Delta W(\beta)^2}{2c} & \text{for } \beta \in (\beta_0,\beta_1). \end{cases}
\end{aligned}$$

From the previous analysis the following facts hold. The first expression is always equal to zero. The expected welfare under the efficient action policy is positive, decreasing and concave in β. The expected welfare under the new action policy is positive, decreasing, initially linear and then convex in β. Moreover, $EW_{NA}(\beta) \to 0$ for $\beta \to \beta_1$. Hence, the SQ regime never (strictly) dominates for any $\beta \in [0,1]$.

Now, let us compare $EW_{EA}(\beta)$ and $EW_{NA}(\beta)$ at $\beta = 0$. We obtain

$$EW_{NA}(0) - EW_{EA}(0) = \frac{\Pi}{c}\overline{W}\frac{\alpha_I^*}{1-\alpha_{II}^*} + g_I(.) + g_{II}(.) > 0,$$

implying that for low probabilities of the bad state β the NA policy is preferred.

Secondly, $EW_{NA}(\beta_1) \to 0$ for $\beta \to \beta_1 < 1$ while $EW_{EA}(1) = 0$. Since both expressions are always decreasing in β with $\frac{\partial^2 EW_{EA}}{\partial\beta^2} < 0$ and $\frac{\partial^2 EW_{NA}}{\partial\beta^2} \geqslant 0$, they intersect only once. Hence, there exists a value $\widetilde{\beta}$ such that $EW_{EA}(\beta) \lesseqgtr EW_{NA}(\beta)$ for $\beta \lesseqgtr \widetilde{\beta}$. ■

Bibliography

Gual, Jordi, Martin Hellwig, Anne Perrot, Michele Polo, Patrick Rey, Klaus Schmidt, and Rune Stenbacka (2005), *An Economic Approach to Article 82*, Report for the DG Competition, European Commission.

Immordino, Giovanni, and Marco Pagano and Michele Polo (2009), "Incentives to innovate and social harm: laissez-faire, authorization or penalties?," CSEF woking paper n.220.

Kaplow, Louis (1994), "The Value of Accuracy in Adjudication: an Economic Analysis," *Journal of legal studies*, **15**, 371-385.

Kaplow, Louis (1995), "A Model of the Optimal Complexity of Legal Rules," *The Journal of Law, Economics and Organization*, **11**, 150-163.

Kaplow, Louis and Steven Shavell (1994), "Accuracy in the Determination of Liability," *Journal of law and economics*, **37**, 1-15.

Kaplow, Louis and Steven Shavell (1996), "Accuracy in the Assessment of Damages," *Journal of law and economics*, **39**, 191-210.

Katsoulakos Yannis and David Ulph, (2009), On Optimal Legal Standards for Competition Policy: a General Welfare-Based Analysis, *Journal of Industrial Economics*, LVII:410-437.

Lear (2006), "The Cost of Inappropriate Intervention and Non Intervention Under Article 82," Report for the Office of Fair Trade, London, UK.

Mookherjee, Dilip and I. Png (1994), "Marginal Deterrence in the Enforcement of Law," *Journal of Political Economy* **102**, 1039-66.

Png I.P.L.(1986), "Optimal Subsidies and Damages in the Presence of Judicial Error," *International Review of Law and Economics*, **6**, 101-105.

Polinsky, Mithcell A. and Steven Shavell (2000), "The Economic Theory of Public Enforcement of Law," *Journal of Economic Literature*, **38**, 45-76.

Posner, Richard A. (1992), *The Economic Analysis of Law*, 4th Edition. Boston: Little Brown.

Schinkel, Maarten P. and Jan Tuinstra (2006), "Imperfect Competition Law Enforcement," *International Journal of Industrial Organization*, **24**, 1267-1297.

Segal Ilya, and Michael Whinston (2007), "Antitrust in Innovative Industries," *American Economic Review*, **97**(5), 1703-1730

Shavell, Steven (1980), "Strict Liability versus Negligence," *The Journal of Legal Studies*, **9**, 1-25.

Shavell, Steven (1992), "A Note on Marginal Deterrence," *International Review of Law and Economics* **12**, 345-355.

Shavell, Steven (2007), "Do Excessive Legal Standards Discourage Desirable Activity?," *Economics Letters*, **95**, 394-397.

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