Journal of Public Economics 95 (2011) 864–876

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Incentives to innovate and social harm: *Laissez-faire*, authorization or penalties?

Giovanni Immordino [a], Marco Pagano [b], Michele Polo [c,*]

[a] Università di Salerno and CSEF, Italy
[b] Università di Napoli Federico II, CSEF, EIEF and CEPR, Italy
[c] Università Bocconi and IGIER, Italy

ARTICLE INFO

Article history:
Received 16 February 2010
Received in revised form 25 January 2011
Accepted 25 January 2011
Available online 2 February 2011

JEL classification:
D73
K21
K42
L51

Keywords:
Innovation
Liability for harm
Safety regulation
Authorization

ABSTRACT

When firms' research can lead to potentially harmful innovations, public intervention may thwart their incentives to undertake research by reducing its expected profitability (average deterrence) and may guide the use of innovation (marginal deterrence). We compare four policy regimes: *laissez faire*, ex-post penalties and two forms of authorization – lenient and strict. If fines are unbounded, *laissez faire* is optimal if the social harm from innovation is sufficiently unlikely; otherwise, regulation should impose increasing penalties as innovation becomes more dangerous. If fines are bounded by limited liability, for intermediate levels of expected social harm it is optimal to adopt (indifferently) penalties or lenient authorization, while strict authorization becomes optimal if social harm is sufficiently likely.

1. Introduction

When private actions generate harmful externalities, public intervention can improve welfare if it appropriately trades off social harm reduction with enforcement costs, as recognized by a vast literature in public economics [1] and in law and economics.[2] Yet, it is rarely recognized that public intervention may stifle innovations that entail benefits as well as risks for society. Even though this idea dates back at least to the work of Friedrich Hayek (1935, 1940), to the best of our knowledge there is no formal analysis of how the design of public policies should take into account the risks and benefits stemming from private innovative activity. [3]

We address this issue, taking into account that public policies may affect both firms' effort to discover new technologies and their actual use, once discovered. Central to our approach is the idea that investment in research and development (R&D) often leads to innovations whose impact on welfare is unknown when the investment is made: not only research may fail to produce workable results, but even if it succeeds, it may lead to innovations with unpredictable welfare effects.

Since generally public policies penalize innovations that turn out to create social hazards, a firm undertaking R&D investment is uncertain as to how public policy will eventually treat the results of its research. Insofar as it expects policy to reduce the expected profitability of innovation, the firm will reduce its R&D investment – a disincentive effect that we label "average deterrence". As we will see, public policies may differ in average deterrence – their research-thwarting effect – as well as in marginal deterrence – their ability to steer innovators towards less harmful implementation of their findings. Precisely these differences dictate which policy is best in each circumstance.

Scientific uncertainty in R&D is an obvious example of the potentially two-edged effects of innovation: research on genetically modified (GM) seeds may pave the way to higher yields in farming, yet pose unknown risks to public health; similar issues arise in the nanotechnology industry (Biello, 2008; Scientific American, 2010) and in the pharmaceutical and chemical industries (Philipson and Sun, 2008). Another example refers to financial innovation: the introduction of new derivatives may open profit opportunities for intermediaries and offer new hedging tools for investors, while creating new dangers for

* Corresponding author. Università Bocconi, Via Sarfatti 25, 20136 Milan, Italy.
E-mail address: michele.polo@unibocconi.it (M. Polo).

[1] Several contributions in public economics highlight that intervention should be curtailed if its enforcement is very costly or generates bribery (Krueger, 1974; Rose-Ackermann, 1978; Banerjee, 1997; Acemoglu and Verdier, 2000; Glaeser and Shleifer, 2003; Immordino and Pagano, 2010, among others). Another strand of research deals with the optimal design of regulation (see Laffont and Tirole (1993) and Armstrong and Sappington (2007)).

[2] This strand of the literature has generated seminal contributions on optimal law enforcement such as Becker (1968), Becker and Stigler (1974) and Polinsky and Shavell (2000).

[3] An exception is the paper by Segal and Whinston (2007) on the impact of antitrust enforcement in high tech industries. Considering a sequence of innovations, the authors analyze the trade-off between protecting the incumbents, that increases the rents of the winner and the incentives to invest in innovation in the first place, and protecting the innovative entrants, that increases the rate of technical progress. They derive conditions under which the latter effect is the dominant one.

0047-2727/$ – see front matter
doi:10.1016/j.jpubeco.2011.01.011

unsophisticated investors who cannot master the information needed to invest in the new securities, as illustrated by the 2007–09 financial crisis. In the words of Lloyd Blankfein, CEO of Goldman Sachs, a key lesson of the crisis is that the financial industry "let the growth in new instruments outstrip the operational capacity to manage them. As a result, operational risk increased dramatically and this had a direct effect on the overall stability of the financial system" (Blankfein, 2009, p. 7).

In each of these situations, society may choose from a range of different regulatory responses. We focus on four different options: (i) *laissez faire*, (ii) a *lenient authorization* regime where inventions can be used commercially if not found to be harmful in tests, (iii) a *strict authorization* regime where they can be used commercially only if ascertained to be beneficial, and (iv) a regime based on *penalties*, where the commercial use of innovations is sanctioned *ex post* if found to be harmful. The difference between authorization and penalty-based regimes does not only lie in the timing of intervention – *ex-ante* scrutiny in the former *versus ex-post* evaluation in the latter – but also in their different degree of flexibility: authorization is a "yes-or-no" decision, and as such it admits no nuances, while penalties can be fine-tuned according to the severity and likelihood of social harm. But even an authorization regime can be designed to be lenient or strict, as just explained, depending on the standards of evidence required about the social effects of innovation.

We show that the greater the social harm that innovations may generate, the more cogent should be the chosen form of public intervention. This general principle applies first of all *within each regime*. When social harm is unlikely a lenient authorization regime is superior to a strict one, while the opposite holds when social harm is likely. Similarly, the penalty regime involves higher fines as the probability of harm increases: in the limiting case of very low risk of social harm, fines are optimally set at zero, effectively leading to a *laissez-faire regime*; as the risk of social harm increases, fines must be gradually increased so as to discourage increasingly harmful actions. This outcome is obtained by setting no fine for actions up to the one the regulator wants to implement, and deterring all other actions by fines large enough as to make them unprofitable. Hence, in the penalty regime, the regulator invariably induces firms to choose the welfare-maximizing action.

This result is no longer feasible if the maximum fine is capped, for instance because of limited liability. Then, penalties cannot deter firms from choosing the actions most harmful to society: these are precisely those yielding the highest profits, so that firms may wish to carry them out even at the risk of paying the maximum fine. In this case, therefore, the penalty regime becomes unappealing if the likelihood of social harm is very high.

The principle that the cogency of public intervention should be increasing in the likelihood of social harm also applies to the choice *across regimes*. If there is no upper bound on fines, society should opt either for *laissez-faire* or for the penalty regime, depending on the likelihood of social harm. In this case, the blunter authorization regimes are invariably dominated. If instead fines are constrained by limited liability, authorization regimes will be preferred when the risk of social harm is sufficiently large. More specifically, in this case the full range of regimes is deployed, depending on the risk level: *laissez-faire* if risk is very low; the penalty regime if risk is in an intermediate range (or equivalently lenient authorization in the top portion of this range); finally, strict authorization for high risk levels. In general, these optimal policies entail underinvestment in research compared to the first-best level, since firms do not internalize the social benefits of innovation (although overinvestment in innovation may occur in the penalty regime when fines are bounded by limited liability).

These policies are softer than those that should be adopted if innovation did not require costly investments in R&D. In that case, regulation would not need to trade off the social risk of social harm with the firm's incentives to innovate, so that only marginal deterrence would matter: *laissez faire* would never be adopted, and in the penalty regime the regulator would deter any harmful action as long as sufficiently high fines are feasible, rather than gradually restricting the firm's choice to less damaging actions as the probability of social harm increases, as done when R&D is costly.

The empirical evidence is consistent with a key prediction of the model – that authorization regimes should be used and become more cogent only when potential social harm is large. In overseeing the safety of medical devices, the FDA authorization process requires more stringent review processes depending on the relevant degree of patient risk.[4] The same principle is now advocated to regulate financial innovation: while the safest securities should be available to investors without authorization, riskier ones, such as derivatives or structured debt, should be sold only upon being authorized, and even so only to eligible investors and in limited amounts.[5] Furthermore, authorization regimes have typically become more stringent when regulators have realized that the likelihood of social harm had been underestimated: the FDA and the European Medicine Agency (EMA) tightened their standards and protocols to authorize drugs since Thalidomide (a morning sickness pill) caused thousands of children in Europe to be born with birth defects in the 1960s.[6]

A second prediction of the model – that tougher regulation comes at the cost of lower incentives to innovate – is also supported by the evidence. The increasingly costly and lengthy clinical trials required by the FDA have prompted growing concerns over the incentives to introduce new drugs: "Ray Hill, president of the British Pharmacological Society … cautions that the much higher costs and larger trials risks reducing pharmaceutical research and stunting innovation. For example, the entire class of Cox 2 painkillers […] was in effect killed by the withdrawal, as the FDA began to demand much bigger pre-approval trials" (Financial Times, 2010b). This outcome would be in line with evidence from the 1960s and 1970s: the annual number of introductions of new chemical entities per dollar of R&D expenditure in the U.S. declined by about sixfold between 1960–61 and 1967–70, while the corresponding figure in the U.K. was threefold. A comparative analysis of these two countries' experience concludes that, controlling for other factors, increased and tighter regulation after 1962 contributed to the slowdown in the innovativeness of the U.S. pharmaceutical industry (Grabowski et al., 1978).

At a theoretical level, our analysis is related to Shavell (1984), who analyzes four determinants of the choice between an authorization and a penalty regime, in his context respectively labeled as safety regulation and liability: (i) difference in risk knowledge; (ii) incentive or ability to enforce penalties; (iii) magnitude of administrative costs, and (iv) magnitude of maximal fines. In our analysis, we hold determinants (i) to (iii) constant across regimes. This is done to focus on the role of innovation in the choice between regimes, eliminating other sources of differential effectiveness between them.

Our model also shares some features with the "activity level" model of law enforcement (Shavell, 1980, 2007; Polinsky and Shavell, 2000). In that model, private benefits and social harm depend on two

[4] The FDA categorizes devices "in one of three classes (I, II, and III), based on the degree of patient risk. Class I devices are the least risky, and typically require no premarket approval from the FDA, although the manufacturer must register with the FDA prior to marketing the device. Class II devices pose more risk to patients, and must receive prior approval via the 501(k) review process, which typically seeks to establish that the given device is substantially equivalent to another device that has received FDA approval. The most risky (class III) devices require approval via the premarket approval process (PMA), which, similar to the process for pharmaceuticals […], involves the submission of a PMA application establishing the device's safety and efficacy, usually through the results of clinical trials" (Philipson at al., 2010, p. 8).

[5] Stephen Cecchetti (Head of the Monetary and Economic Department of the BIS) argues that, just like drugs must undergo clinical testing before being authorized for sale, financial offerings would be subject to similar tests before being authorized: "An instrument could move to a higher category of safety only after successful tests analogous to clinical trials" (Financial Times, 2010a).

[6] "Thalidomide marked a turning point in the history of drug regulation, leading the authorities around the world to impose higher approval standards to insure drugs were tested for safety as well as efficacy" (Financial Times, 2010b).

different decisions by agents – an activity level (say, how long an individual drives a car) and a level of precaution (driving speed) – and the analysis typically compares the effects of different liability rules (strict versus fault-based liability). Our innovative activity is reminiscent of the activity level, while the choice of new actions parallels the choice of precaution. But our timing and information structure differ from those of the standard activity model. There, agents typically choose activity and precaution simultaneously and perfectly know the effect of their actions on welfare and the rules that will apply to them; the design of these rules aims at steering their choices so as to minimize social harm. So the issue is only one of marginal deterrence. In contrast, in our model when firms choose their research effort, they still ignore whether it will produce a beneficial or a harmful innovation, and therefore consider the policies designed for both cases as potentially relevant to them. Due to this veil of ignorance, policies devised to penalize socially harmful innovations may end up deterring research by firms that would in fact produce beneficial innovations. That is why uncertainty is key to what we call average deterrence.

The model that comes closest to ours is that of Schwartzstein and Shleifer (2009), who investigate when and how the optimal policy combines ex-ante regulation and ex-post litigation in the activity model. They consider a setting where safe and unsafe firms decide whether to produce and may take precautions. Firms face uncertainty as to the liability for damages that will apply to them, due to possible judicial errors: a judge may mistake a safe firm for an unsafe one, which creates a disincentive effect for safe firms. If the regulator can identify safe firms *ex ante* (or at least can do so better than courts), it is optimal for regulation to set these firms free from liability for damages, since the social benefits of their activity exceeds the expected harm from taking too few precautions. This parallels our finding that regulation should be softer when social harm is unlikely. But our analysis differs in the way uncertainty is modeled: in our setting, it is an inherent feature of firms' research activity, rather than an effect of judicial errors. As such, it applies uniformly to any form of policy intervention, and does not *per se* favor any regime over others.

The paper is organized as follows. Section 2 presents the model. Section 3 presents three benchmark cases: the first best, where the regulator directly controls firms' choices, *laissez faire*, where firms are unrestricted, and per-se illegality, where the new actions are always prohibited. Section 4 analyzes the authorization regime, Section 5 the penalties regime, and Section 6 the overall optimal policy. Section 7 discusses several extensions, and Section 8 concludes. The proofs are in the Appendix A.

2. Setup

We consider a profit-maximizing firm that must choose whether to invest in R&D activity or not. If the firm does not invest in such activity, it can select only among known actions, e.g. familiar technologies. If instead the firm invests and succeeds in its research effort, it expands its opportunity set. However, the new actions made possible by innovation, though expected to be profitable, may have *ex-ante* unknown social effects. For instance, a biotech firm may produce traditional seeds or experiment with new GM seeds that promise higher yields but pose unknown risks to public health.

To contain the potential hazards posed by innovative activity, public policy may constrain the actions of successful innovators either by subjecting them to an *ex-ante* notification and authorization requirement (*authorization*) or to an *ex-post* penalty regime (*penalty*). Under the authorization regime, the firm notifies to a public agency (such as the Food and Drugs Administration) the action it plans to undertake based on the results of its research (e.g., the sale of GM seeds), and the agency decides whether the firm is allowed to go ahead, after carrying out an investigation on the potential implied harm. In contrast, under a penalty regime the firm is free to choose any new action made possible by its research findings (in our example, sell any new GM seed), but may have to pay a fine *ex-post* if this action causes social harm. Public policies must trade off the social gains arising from the firm's innovation (a larger harvest) against their potential social harm (a public health hazard). The key issue to be explored is how this trade-off shapes the optimal design of policy in each regime, as well as the choice between regimes. To focus the analysis on the role of innovation in the choice between regimes, we neglect other sources of differential effectiveness between them, by assuming that the probability p of finding evidence about the social effects of the innovation is the same across regimes, irrespective of whether it is collected *ex ante* in an authorization procedure or *ex post* in assessing the social harm caused by the innovation.

In our analysis, the firm is assumed to know how to implement the *status-quo* action a_0 (selling traditional seeds), as well as the associated profits Π_0 and welfare W_0, which are normalized to zero with no loss of generality: $\Pi_0 = W_0 = 0$. In contrast, carrying out a new action requires innovative activity (experiments with GM seeds). If the investment is unsuccessful, the firm must implement the *status-quo* action a_0. If it is successful, the firm discovers how to implement a set of new actions $A = (0, \bar{a}]$, with associated profits $\Pi = \pi a$, where $\pi > 0$.[7] In this case, the firm is also assumed to learn the state of nature $s \in \{b, g\}$: in the bad state b, the innovation is socially harmful, whereas in the good state g it is beneficial. Proceeding with our example, the biotech company learns not only how to produce new GM seeds, but also the dangers that they pose to public health.

Depending on the state of nature s, the social consequences of new actions are described by one of two different functions. In state b, which occurs with probability β, new actions decrease welfare according to $W^b = -w^b a$, with $w^b > 0$. In the bad state, private incentives conflict with social welfare since a new action a yields profit πa but reduces welfare by $w^b a$. Hence, the probability β measures the misalignment between public interest and firms' objectives: in our example, β is the prior probability that GM seeds will pose a health hazard. Instead, in the good state g, that occurs with probability $1 - \beta$, new actions raise welfare according to the function $W^g = w^g a$. In this state, the social gains from innovation exceed private ones, that is, $w^g > \pi$, or equivalently new actions increase consumer as well as producer surplus.[8]

The resources I that the firm invests in research determines its chances of success: for simplicity, I is assumed to coincide with the success probability, so that $I \in [0, 1]$. The cost of learning is increasing and convex in the firm's investment. For concreteness we assume

$$c(I) = c\frac{I^2}{2},$$

where $c > w^g \bar{a}$ ensures an internal solution lower than 1 for the choice of I in all the regimes that we shall consider, including the first best. After choosing its investment in research I and learning its outcome, the firm selects the most profitable action among the feasible ones (which include the *status-quo* action) under the constraints imposed by public policy.

[7] The assumption that the new action is always profitable may be taken as reflecting a regime of patent protection for innovations. However, in the model we assume that patents are released by a different agency based on objectives and criteria independent of those of the regulator considered in the model. Therefore, patent release criteria do not belong to the set of tools considered in the regulatory design problem of this paper.

[8] The assumption that profits and welfare are linear in actions allows us to compare analytically the authorization and liability regime. Without linearity, this comparison could only be effected via numerical simulations. However, as shown in Section 7.1 below, this assumption is not equivalent to postulating risk-neutrality of social preferences.

3. Benchmarks: first best, *laissez faire* and *per-se* illegality

As the opportunities created by innovation generate positive or negative externalities, depending on the state of nature, public policy may be beneficial. To evaluate public interventions, it is useful to compare them against three benchmarks: (i) the first-best outcome (*FB*), which would obtain if the regulator could control firms' choices I and a directly; (ii) a *laissez faire* regime (*LF*), where firms are free to choose whichever action they like; and (iii) a *per-se* illegality regime (*PI*), where investing in research is forbidden altogether.

We define a_r^g and a_r^b, respectively, the action taken in the good and bad state under the policy regime r. Unconstrained welfare maximization calls for action $a_{FB}^g = \bar{a}$ in the good state and action $a_{FB}^b = a_0 = 0$ in the bad state, so that the first-best expected welfare is

$$E(W_{FB}) = I\left[(1-\beta)w^g a_{FB}^g - \beta w^b a_{FB}^b\right] - c\frac{I^2}{2} = I(1-\beta)w^g\bar{a} - c\frac{I^2}{2}. \quad (1)$$

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

$$I_{FB} = \frac{(1-\beta)w^g\bar{a}}{c}, \quad (2)$$

which is increasing in the likelihood of the good state $1-\beta$ and in the associated welfare gain $w^g\bar{a}$, and decreasing in the marginal cost of innovative activity c.

In the polar opposite scenario of *laissez faire*, firms maximize profits without constraint from public policies: a firm will opt for the most profitable action whenever its research has been successful, irrespective of the state, i.e. $a_{LF}^g = a_{LF}^b = \bar{a}$. Its expected profits from innovation are $E(\Pi) = I\pi\bar{a} - cI^2/2$, which are maximized if investment in research is

$$I_{LF} = \frac{\pi\bar{a}}{c}, \quad (3)$$

which may exceed the first-best level in Eq. (2) when the bad state is very likely or fall short of it when β is low. The welfare level associated with the *laissez-faire* level of investment is

$$E(W_{LF}) = I_{LF}\left[(1-\beta)w^g\bar{a} - \beta w^b\bar{a}\right] - c\frac{I_{LF}^2}{2} = \frac{\pi\bar{a}^2}{2c}(2E(w) - \pi), \quad (4)$$

where we denote the expected marginal welfare of action $\bar{a}$ by $E(w) \equiv (1-\beta)w^g - \beta w^b$.

Finally, if public policy makes innovative activities *per-se* illegal, then welfare is trivially $E(W_{PI}) = 0$.

In the following analysis, the first best will be unattainable, because the policy maker is assumed not to control firms' choices directly, but to influence them either via authorizations or via penalties. We assume policy makers to be benevolent, in the sense that they design and enforce policies so as to maximize social welfare. Since under any regime public decisions are taken according to this goal, we can avoid defining precisely the institutional framework in which the public policies are designed and enforced. Henceforth we just refer to an "agency", which might be a legislator, a regulator, an authority or a judge depending on the relevant regime.

4. Authorization

In the authorization regime, after a firm notifies the action that it wishes to undertake, the authorizing agency investigates whether the notified actions are socially harmful or not, and obtains decisive evidence about their social effects with probability $p \in [0,1]$, while it finds no evidence in either direction with probability $1-p$. If the evidence is decisive, the authorization is given if and only if the evidence is favorable. If the evidence is not decisive, instead, the agency can opt for one of two rules: a "lenient authorization" (*LA*) rule whereby when in doubt the firm is authorized, or a "strict authorization" rule (*SA*) whereby in such circumstances the authorization is denied. Hence, under the *LA* regime the firm is authorized as long as no social harm is proved, while under the *SA* rule new actions are permitted only if proved to be socially beneficial. The two regimes, therefore, differ if the enforcer's evidence is not decisive, in which case the firm is authorized in the *LA* regime, while it is not in the *SA* regime.

If the preliminary review were always to produce decisive evidence ($p = 1$), the two regimes would be equivalent; but if it may fail to yield hard evidence ($p < 1$), the lenient and strict rules differ. *LA* leads to under-enforcement, since with probability $\beta(1-p)$ it gives green light to a harmful action, while *SA* entails over-enforcement, by blocking with probability $(1-\beta)(1-p)$ a beneficial action. When the authorization is denied, the firm must take the *status-quo* action a_0. In principle, the agency may also opt for "*per-se* illegality", by always denying authorization. But this option is invariably dominated by a strict authorization regime, as we shall see below.

In the authorization regime, the timing of the game is as follows. At $t = 0$ the agency chooses between the *LA* and the *SA* regime, committing to the chosen rule for the entire game. At $t = 1$ the firm chooses its innovative activity I and with probability I discovers the new actions A and the state of nature s. At $t = 2$, in regimes *LA* and *SA* the firm notifies the agency of the new action it wishes to undertake. At $t = 3$ the agency obtains evidence on the social effects of the proposed action with probability p, and decides whether to authorize it or not. At $t = 4$ the firm carries out the authorized action (if any), and the corresponding private and social payoffs are realized.

Since by assumption the new actions in A are more profitable than the *status-quo* action a_0, if research is successful the firm always applies to be authorized to carry out the highest (most profitable) new action $\bar{a}$.[9] In the *LA* regime, the firm anticipates that the agency will always authorize it in the good state (whether it uncovers favorable evidence or not), i.e. $a_{LA}^g = \bar{a}$ and will authorize it only with probability $1-p$ in the bad state (that is, only if no decisive evidence is uncovered). Hence the expected action in the bad state will be $a_{LA}^b = pa_0 + (1-p)\bar{a} = (1-p)\bar{a}$. In the *LA* regime the firm therefore will take action $\bar{a}$ with probability $(1-\beta) + \beta(1-p) = 1-\beta p$, and its expected profits when the investment is chosen are

$$E(\Pi_{LA}) = I(1-\beta p)\pi\bar{a} - c\frac{I^2}{2},$$

so that its optimal innovative activity is

$$I_{LA} = \frac{(1-p\beta)\pi\bar{a}}{c}. \quad (5)$$

Notice that, since when I is chosen the firm does not yet observe whether the innovation will be socially beneficial or harmful, it takes into account how public policy treats both occurrences. This feature applies to the authorization regimes as well as to the penalty regime to be discussed in the next section. The expected welfare under lenient authorization is

$$\begin{aligned} E(W_{LA}) &= I_{LA}\left[(1-\beta)w^g\bar{a} - \beta(1-p)w^b\bar{a}\right] - c\frac{I_{LA}^2}{2} \\ &= \frac{\pi\bar{a}^2}{2c}(1-p\beta)\left[2E(w) - \pi + p\beta\left(2w^b + \pi\right)\right] \end{aligned} \quad (6)$$

where, as above, $E(w) \equiv (1-\beta)w^g - \beta w^b$ is the marginal social value of the new actions.

[9] Since the authority obtains the same (decisive or null) evidence on all the new actions analyzed, it has to apply the same response to any of them. Hence, it is equivalent for the firm to require an authorization on all the new actions A and then pick up $\bar{a}$ or just for the selected action $\bar{a}$.

Under the *SA* regime, instead, the agency will authorize action $\bar{a}$ only if it uncovers favorable evidence, which happens only in the good state. Hence $a^g_{SA} = p\bar{a} + (1-p)a_0 = p\bar{a}$ and $a^b_{SA} = a_0$. Since action $\bar{a}$ will be authorized with probability $(1-\beta)p$, the firm's expected profits when the investment is chosen are

$$E(\Pi_{SA}) = I(1-\beta)p\pi\bar{a} - c\frac{I^2}{2},$$

so that its optimal innovative activity is

$$I_{SA} = \frac{(1-\beta)p\pi\bar{a}}{c}. \tag{7}$$

Clearly, the lenient rule is associated with greater investment in innovation than the strict one ($I_{LA} > I_{SA}$), because it leaves greater expected profits to innovators. The welfare level associated with the *SA* regime is

$$E(W_{SA}) = I_{SA}(1-\beta)pw^g\bar{a} - c\frac{I_{SA}^2}{2} = \frac{\pi\bar{a}^2}{2c}(1-\beta)^2p^2(2w^g - \pi). \tag{8}$$

The following lemma establishes that the lenient rule – being more permissive towards innovators – is optimal if and only if innovation is sufficiently unlikely to cause social harm:

Lemma 1. Optimal authorization

When $p = 1$ the lenient and strict authorization regimes are equivalent, while if $p < 1$ there exists a value $\beta_{LA}(p) \in [0,1]$ such that the LA regime is weakly preferred to the SA one iff $\beta \leqslant \beta_{LA}(p)$. The threshold $\beta_{LA}(p)$ is increasing in p and tends to 1 as $p \to 1$.

Intuitively, the *LA* regime, being associated to under-deterrence, is more favorable to boost innovative investment than the *SA* regime. When there is a low probability that the new actions reduce welfare, the former is preferable to the latter. This happens for a larger set of values of the probability β when enforcement becomes more effective (higher p): when the probability of social harm increases, the agency sticks to the lenient regime only if the ability to detect harmful innovations is high enough as to compensate the under-deterrence of this regime. Finally, notice that the strict authorization regime weakly dominates *per se* illegality, since it allows to implement the new action when it is sure to increase welfare, while *per se* illegality forgoes this opportunity.

5. Penalties

In the penalty regime, to be denoted by *P*, successful innovators can implement their preferred action $\bar{a}$ but anticipate that they may be charged a fine if the action is found to have caused social harm. This occurs when the agency obtains definite evidence that the chosen action was socially harmful, which occurs with probability p as in the authorization regime, as already mentioned in Section 2. In this regime, an action $a \in A$ that causes social harm relative to the *status quo* $(-w^b a < 0)$ is punished according to a fine schedule $f(a)$ chosen in the interval $[0, F]$ and non-decreasing in social harm. [10] This legal rule, that in our example would prohibit to sell hazardous GM seeds, is effect-based, as it punishes only actions that are *ex-post* socially damaging and does so in proportion to the harm caused. We do not set any upper bound on the maximum fine F at this stage, although in Section 6 we shall consider how the optimal policy changes when limited liability constrains the fine schedule, namely $F = \pi\bar{a}$.

In the penalty regime, the timing of the game is as follows: at $t = 0$ the agency commits to the penalty regime and to the fine schedule $f(a)$. If the agency sets its fines at zero for any new action $a \in A$, it effectively opts for the *laissez-faire* regime. At $t = 1$, the firm chooses innovative activity I and with probability I discovers the set of new actions A and the state of nature s. At $t = 2$ it decides which action a to take. At $t = 3$ the private and social payoffs are realized. At $t = 4$ the agency investigates the action a, finds decisive evidence about its social effects with probability p and, if it does, levies the fine $f(a)$.

The choice of actions at $t = 2$ depends on the outcome of the firm's innovative activity at $t = 1$ and on the fine schedule $f(a)$ designed by the agency at $t = 0$. When innovative activity is unsuccessful, the firm carries out the *status-quo* action a_0. Instead, when successful the firm can also take new actions $a \in A$. If these are socially beneficial, all of them are lawful, so that the firm picks the most profitable action, i.e. $a^g_P = \bar{a}$. If instead the new actions $a \in A$ are socially harmful, they are illegal and, if chosen, are sanctioned by a fine. The firm chooses the action that maximizes its profits, net of the expected fine:

$$a^b_P = \tilde{a} = \arg\max_{a \in \{a_0, A\}} [\pi a - pf(a)]. \tag{9}$$

Notice that the firm can always opt for the *status-quo* action that yields zero profits, so that $\pi\tilde{a} - pf(\tilde{a}) \geq 0$. Referring again to our example, if innovative activity is unsuccessful, the firm sells traditional seeds, while if successful it markets the most profitable type of seeds if that poses no concern for public health, while it selects a less profitable variety if it is dangerous, taking into account the corresponding fines it may be called to pay. We summarize this discussion as follows:

Lemma 2. Actions

At stage 2, given the fine schedule $f(a)$, the firm chooses (i) a_0 if its innovative activity is unsuccessful; (ii) $a^g_P = \bar{a}$ if it is successful and the new actions are socially beneficial; (iii) $a^b_P = \tilde{a}$ as defined by (9) *if it is successful and the new actions are socially harmful.*

At stage 1 the firm chooses the innovative activity I so as to maximize its expected profits, anticipating the optimal actions to be taken at stage 2. In terms of our example, the biotech firm chooses its investment in R&D, taking into account which GM seeds it will sell if successful. Its expected profits at this stage are:

$$E(\Pi_P) = I[(1-\beta)\pi\bar{a} + \beta(\pi\tilde{a} - pf(\tilde{a}))] - c\frac{I^2}{2}. \tag{10}$$

The expression in square brackets is the expected gain from innovative activity, net of expected fines. This expression is always positive since, as argued above, $\pi\tilde{a} - pf(\tilde{a}) \geq \Pi_0 = 0$. Hence, the firm will always perform some innovative activity. As already observed under the authorization rules, also in the penalty regime the expected profits and the return to innovative investment depend on the policies applied in the good and bad states, since when the investment is chosen the welfare effect of the new actions is still unknown.

Maximizing Eq. (10) with respect to I yields:

Lemma 3. Innovative activity

At stage 1, given the fine schedule $f(a)$, the optimal private level of innovative activity is

$$I_P(\tilde{a}, f(\tilde{a})) = \frac{(1-\beta)\pi\bar{a} + \beta[\pi\tilde{a} - pf(\tilde{a})]}{c}. \tag{11}$$

[10] We do not consider negative fines, i.e. subsidies to innovation, since – if present – these are generally chosen by authorities that are not the same as those in charge of preventing the socially harmful effects of firms' innovative activities. However, if any other branch of government provides subsidies to innovation, these will be captured by an increase in our profitability parameter π: the presence of subsidies to innovation will lead policy to shift towards stricter rules, as shown below in Section 7.3.

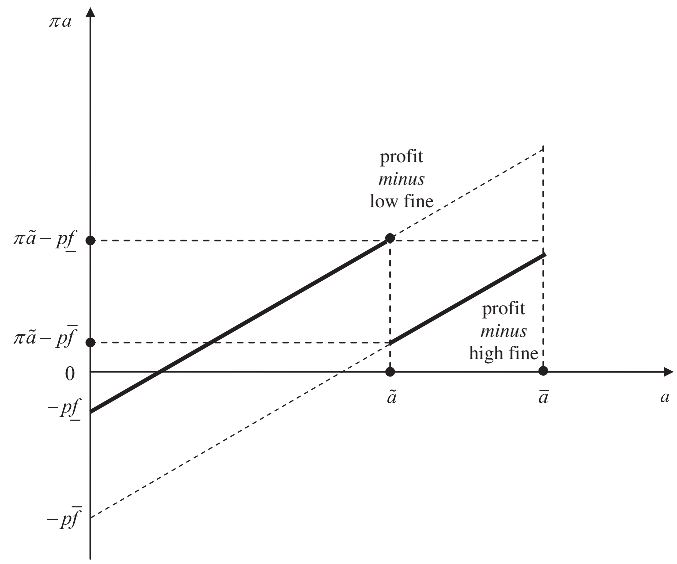

Fig. 1. Actions, profits and fines.

We now turn to the design of the fine schedule at stage 0. The influence of law enforcement and penalties on firms' behavior is twofold: it affects both the choice of the action $\tilde{a}$ when innovation succeeds and new actions are unlawful, and the incentives to pursue innovative activity I in the first place, when the firm computes the expected profits from the new actions taking into account that in the bad state it will get $\pi\tilde{a} - pf(\tilde{a})$. The first role of law enforcement is known in the literature on law enforcement as *marginal deterrence*, that is, the ability of fines to guide private choices among unlawful actions. [11] The second role, which is absent in standard models, stems from the impact of law enforcement on innovative activity, and therefore on the probability that any new action a will be taken. For this reason we label this second effect *average deterrence*. The policy parameters will be chosen considering both effects on private choices and ultimately on welfare.

A convenient way to describe the optimal penalty schedule is to consider its design as equivalent to (indirectly) implement a (profit-maximizing) action $\tilde{a}$, that is an action that the firm will want to pick in the bad state, and to select among all these implementable actions $\tilde{a}$ the one that maximizes the expected welfare, which we denote as $\hat{a}$. Since the profit function πa is increasing in a, if we want to implement action $\tilde{a}$ we can use with no loss of generality, within the set of non-decreasing fine schedules, the stepwise function

$$f(a) = \begin{cases} \underline{f} \geqslant 0 & \text{if } a \leqslant \tilde{a} \\ \overline{f} \leqslant F & \text{if } a > \tilde{a} \end{cases} \tag{12}$$

where we choose $\tilde{a}$, $\underline{f}$ and $\overline{f}$ in order to satisfy the incentive compatibility constraint

$$\pi\tilde{a} - p\underline{f} \geqslant \pi\overline{a} - p\overline{f} \tag{13}$$

or equivalently

$$\overline{f} \geqslant \underline{f} + \frac{\pi(\overline{a} - \tilde{a})}{p}.$$

We rely on Fig. 1 to illustrate this point. The function in Eq. (12) shifts the profit function πa downward by $p\underline{f}$ to the left of point $\tilde{a}$, and by $p\overline{f} > p\underline{f}$ to its right. By appropriately choosing the fines $\underline{f}$ and $\overline{f}$, the regulator can turn $\tilde{a}$ into the firm's global optimum, thus satisfying the condition in Eq. (13). Then, the firm will choose $\tilde{a}$ and earn profits $\pi\tilde{a} - p\underline{f}$ in the bad state. Among all the implementable actions, the public policy will choose $\underline{f}, \overline{f}$ and $\tilde{a}$ so as to maximize expected welfare,

$$\mathrm{E}(W_P) = I_P\left(\tilde{a}, \underline{f}\right)\left[(1-\beta)w^g\overline{a} - \beta w^b\tilde{a}\right] - c\frac{I_P^2\left(\tilde{a}, \underline{f}\right)}{2}, \tag{14}$$

subject to the incentive compatibility constraint, which can be conveniently rewritten as

$$\tilde{a} \geqslant \overline{a} - p\frac{\overline{f} - \underline{f}}{\pi}. \tag{15}$$

Notice that fines do not enter in the expression of the expected welfare, being pure transfers. This maximization program is solved by the following first-order conditions:

$$\begin{aligned} \frac{\partial \mathrm{E}(W_P)}{\partial \tilde{a}} &= [\mathrm{E}(W) - cI_P]\frac{\beta\pi}{c} - \beta w^b I_P + \lambda \geqslant 0, \\ \frac{\partial \mathrm{E}(W_P)}{\partial \underline{f}} &= -[\mathrm{E}(W) - cI_P]\frac{\beta p}{c} - \lambda\frac{p}{\pi} \leqslant 0, \\ \frac{\partial \mathrm{E}(W_P)}{\partial \overline{f}} &= \lambda\frac{p}{\pi} \geqslant 0, \end{aligned} \tag{16}$$

where $\mathrm{E}(W) = (1-\beta)w^g\overline{a} - \beta w^b\tilde{a}(\cdot)$ is the expected welfare from the innovative activity and the term in squared brackets is its expected marginal social value. Finally, the complementary slackness condition is

$$\lambda\left(\tilde{a} - \overline{a} + p\frac{\overline{f} - \underline{f}}{\pi}\right) = 0. \tag{17}$$

The partial derivative $\partial E(W_P)/\partial\tilde{a}$ measures the welfare effect of allowing the firm to choose a more profitable (though more harmful) action $\tilde{a}$ in the bad state, that is, to reduce deterrence. The expression reveals that reduced deterrence affects welfare via two channels. An indirect channel (corresponding to the first term) acts through greater incentives to innovate, i.e. through lower average deterrence: $\mathrm{E}(W) - cI_P$ measures the net social gain to investment in research $(\partial E(W_P)/\partial I_P)$, whereas $\beta\pi/c$ is the effect of allowing a more profitable action $\tilde{a}$ on such investment $(\partial I_P/\partial\tilde{a})$. A direct channel (the second term) instead operates by letting successful innovators choose a more profitable (though more harmful) action $\tilde{a}$ in the bad state, i.e. through lower marginal deterrence. Interestingly, the first term is positive (as shown in the proof of Lemma 4), while the second is negative: allowing firms to go for more profitable actions raises welfare by fostering innovation, but reduces it insofar as it lets successful innovators choose more harmful actions in the bad state. In the optimal policy, the implementable action $\tilde{a}$ is set at the level $\hat{a}$ that maximizes welfare. The first effect (average deterrence) plays an important role in shaping the optimal policy when $w^g - w^b - \pi > 0$, which clearly holds when the positive welfare effect of innovation in the good state is sufficiently large: under this condition the agency will adopt a *laissez-faire* policy if the probability of the bad state is sufficiently low, whereas he will adopt an optimal penalty policy if this probability is high. More precisely:

Lemma 4. *Optimal penalties* If $w^g - w^b - \pi > 0$ *the optimal policy depends on the probability of the bad state* β:

(i) *for* $\beta \in [0, \beta_{LF}]$, *where*

$$\beta_{LF} = \frac{w^g - w^b - \pi}{w^g + w^b}, \tag{18}$$

the optimal policy is laissez faire: it implements $\hat{a} = \overline{a}$, *by setting* $\underline{f} = 0$ *and any* $\overline{f} \geq 0$;

[11] See the seminal work by Stigler (1970) and, for a more general treatment, Mookherjee and Png (1994).

(ii) *for* $\beta \in [\beta_{LF}, 1]$, the optimal policy is a penalty regime that *implements the action*

$$\hat{a}(\beta) = \bar{a}\frac{(1-\beta)}{\beta}\frac{\left(w^g - w^b - \pi\right)}{(2w^b + \pi)}, \quad (19)$$

by setting $\underline{f} = 0$ *and*

$$\bar{f} \geq \frac{\pi\bar{a}}{p}\frac{\beta\left(w^g + w^b\right) - \left(w^g - w^b - \pi\right)}{\beta(2w^b + \pi)}. \quad (20)$$

If $w^g - w^b - \pi \leq 0$ the optimal policy for any β is a penalty regime that implements $\hat{a}(\beta) = 0$.

This lemma implies that, when innovation may have large social benefits and social harm is sufficiently unlikely ($w^g - w^b - \pi > 0$ and $\beta \in [0, \beta_{LF}]$), the firm should be allowed to choose the most profitable action $\bar{a}$ even in the bad state, in order to foster investment in innovation: in this case the penalty regime is effectively equivalent to the *laissez-faire* regime and encompasses it, as witnessed by the fact that the implemented action is the one preferred by the firm. Notice that the *laissez-faire* interval shrinks if the marginal social loss w^b or the marginal profitability π increase, since both these parameter shifts move the boundary β_{LF} to the left.

When social harm is sufficiently likely, i.e. $\beta \in [\beta_{LF}, 1]$, the firm must be constrained by a fine $\bar{f}$ large enough as to implement the action $\hat{a}(\beta) < \bar{a}$ in the bad state. The implemented action $\hat{a}(\beta)$ varies continuously from the most profitable one ($\bar{a}$) to the *status-quo* action ($a_0 = 0$) as the probability of the bad state β, as well as the marginal social loss w^b, increase. The new actions are completely deterred only in the limiting case $\beta = 1$, while for $\beta < 1$ some welfare-decreasing actions $\hat{a}(\beta) > 0$ are accepted. Moreover, the low fine is invariably set at zero ($\underline{f} = 0$), in order to sustain innovative investment.

When instead the innovation produces low social benefits ($w^g - w^b - \pi \leq 0$), for instance because the marginal social loss w^b is extremely large, the public agency will be less forgiving: *laissez faire* is never an option. Quite to the opposite, complete deterrence obtains: in the bad state the *status-quo* action is implemented for any β.

It is important to note that the optimal policy is always able to implement the welfare-maximizing action $\hat{a}(\beta)$ in the interval $[0, \bar{a}]$, since the maximum fine is unbounded: even a poorly profitable action close to 0 is implementable, if needed to maximize welfare, since it is possible to set the maximum fine at a sufficiently high level to discourage the firm from switching to the most profitable action $\bar{a}$ and pay the high fine $\bar{f}$.

Replacing the optimal choices Eqs. (11), (19) and $\underline{f} = 0$ in Eq. (14) yields:

$$\mathrm{E}(W_P) = \frac{\pi\bar{a}^2}{2c}(1-\beta)^2\frac{\left(w^g + w^b\right)^2}{(2w^b + \pi)} = \frac{\pi\bar{a}^2}{2c}\frac{(1-\beta)^2}{(1-\beta_{LF})^2}\left(2w^b + \pi\right). \quad (21)$$

6. Optimal choice of regime

We are now equipped to derive the optimal choice of regime, by comparing the expected welfare associated with each of them:

Proposition 1. *Optimal policy If there is no upper bound on feasible fines and* $w^g - w^b - \pi > 0$, *the optimal (second-best) policy is to adopt laissez faire for* $0 \leq \beta \leq \beta_{LF}$ *and the penalty regime for higher* β. *If instead* $w^g - w^b - \pi \leq 0$ *the optimal (second-best) policy entails adopting the penalty regime for any* β.

According to our result, public intervention becomes increasingly stringent as the danger of social harm increases: as β goes up, the optimal policy changes from *laissez-faire* to a penalty regime. The *laissez-faire* regime is dominated only if the marginal social loss w^b is very large, in which case positive fines are given for actions exceeding $\hat{a}$. Instead, the authorization regime is always dominated. This is because penalties can be fine-tuned to the likelihood of social harm, whereas authorizations are more rigid, being "yes-or-no" decisions that do not affect at the margin the choice of the action. The penalty regime implicitly tolerates, by setting a zero fine, a welfare-decreasing action $\hat{a}(\beta) > 0$ in the bad state, in order to sustain innovative investment. In this sense, the penalty regime entails some form of under-deterrence, as the lenient authorization regime, which however is less subtle than the penalty regime in its deterrence effects.

With unbounded fines, the penalty regime dominates since it is always able to implement the welfare maximizing action $\hat{a}(\beta)$, no matter how low it is, by setting a sufficiently high maximum fine $\bar{f}$. It is therefore interesting to analyze the optimal policies in an alternative environment where the maximum feasible fine F is capped at some upper bound. A natural choice for this bound is the limited liability rule, that constrains the maximum fine not to exceed the firm's maximum profits, that is, $F = \pi\bar{a}$. This limited liability constraint makes the optimal policy considerably richer: while in its absence only *laissez faire* and the penalty regime are used (for low and high values of β, respectively), in its presence also the authorization regimes play a role for sufficiently large values of the likelihood of social harm β. More specifically, the limited liability rule hampers the effectiveness of the penalty regime for values of β exceeding the threshold:

$$\beta_P(p) = \frac{w^g - w^b - \pi}{w^g + w^b - p(2w^b + \pi)}. \quad (22)$$

For an interval of values above this threshold, the penalty and the lenient authorization regimes become equivalent, while for even larger values of β the strict authorization regime dominates both of them. Intuitively, when social harm is very likely and fines are capped at a maximum, the incomplete deterrence of the penalty regime becomes too costly for society. At that point, the strict authorization regime dominates, being safer though less sophisticated. Formally:

Proposition 2. Optimal policy with limited liability

If feasible fines cannot exceed the firm's profits, i.e. $F = \pi\bar{a}$ *and* $w^g - w^b - \pi > 0$, *the optimal (third best) policy entails laissez faire for* $\beta \in [0, \beta_{LF}]$, *penalties for* $\beta \in (\beta_{LF}, \beta_P(p)]$, *penalties or lenient authorization for* $\beta \in (\beta_P(p), \beta_{LA}(p)]$ *and strict authorization for* $\beta \in (\beta_{LA}(p), 1]$. *If instead* $w^g - w^b - \pi \leq 0$, *only lenient authorization up to* $\beta_{LA}(p)$ *and then strict authorization are chosen.*

The threshold $\beta_P(p)$ in Eq. (22) is increasing in p, as shown in Fig. 2, when $w^g - w^b - \pi > 0$[12]: the greater the effectiveness of enforcement, as measured by the probability of detection p, the greater the value of β for which the cap on the maximum fine becomes binding, and therefore the larger the interval where the penalty regime dominates lenient authorization. This is captured in Fig. 2 by the fact that the penalty regime area widens as p increases. In the limiting case of perfect enforcement, i.e. $p = 1$, penalties always dominate authorization, as indicated by the fact that in this case Eq. (22) yields $\beta_P(p) = 1$. Once again, if the marginal social loss w^b is very large the policy becomes more rigid and only the authorization regimes are chosen.

So far, our analysis has disregarded the enforcement costs of the policies. However, the above results would be qualitatively unchanged if the information needed for enforcement has a fixed collection cost, consistently with our assumption that such information collection

[12] If $w^g - w^b - \pi = 0$, then $\beta_P(p) = 0$ for any p. If instead $w^g - w^b - \pi < 0$, then $\beta_P(p) < 0$ for any p.

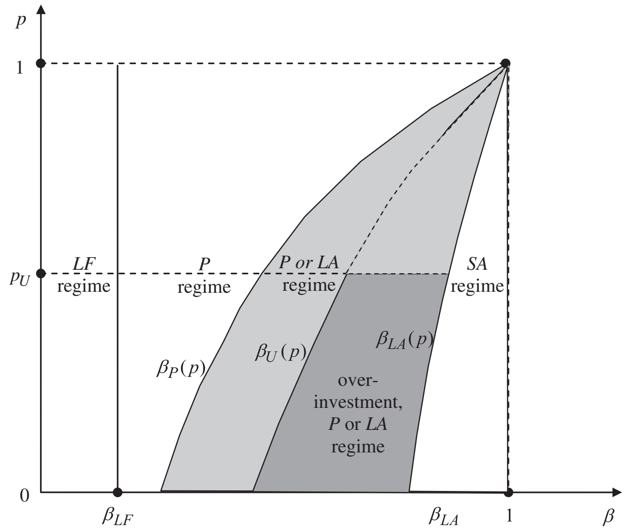

Fig. 2. Investment in innovation and optimal policy with limited liability.

yields decisive evidence with a constant probability p.[13] The presence of enforcement costs would simply expand the region where *laissez faire* dominates (as this regime entails no enforcement costs) at the expense of the penalty regime: β_{LF} would shift to the right, leaving the other boundaries unaffected. If the fixed cost of investigation were to differ between the penalty and authorization regimes, due to different standards or technologies of proof, the thresholds would shift accordingly: for instance, if the penalty regime featured larger enforcement costs than the authorization regime, penalties would be preferred to authorization over a narrower range of parameters.

Having identified the optimal policies for different values of the likelihood of social harm β and detection probability p, it is worth considering whether these policies are associated with under- or overinvestment in innovation by firms:

Proposition 3. Innovative investment

If there is no upper bound on feasible fines, the optimal policy involves underinvestment. If instead feasible fines cannot exceed the firm's profits, the optimal policy entails overinvestment in a region in the (p,β) space and underinvestment elsewhere.

The first part of this proposition refers to the case where there are no constraints on fines in the penalty regime. In this case, underinvestment occurs because investment in R&D is chosen by firms so as to maximize their profits rather than social welfare; moreover, in the penalty regime such investment is discouraged by fines.

The second part of the proposition, which refers to the case where fines are constrained by limited liability, is best understood by looking at Fig. 2. As shown in the proof of the proposition, overinvestment only occurs in the dark shaded area comprised between the two increasing functions $\beta_U(p)$ and $\beta_{LA}(p)$ and for detection probabilities p below a threshold value p_U. In this area, the penalty regime and the lenient authorization regime are both optimal, and at the right boundary $\beta_{LA}(p)$ both become dominated by the strict authorization regime. Elsewhere in the diagram, there is always underinvestment, as in the previous case. Intuitively, in this region firms overinvest in innovation because here the limited liability constraint binds, and therefore blunts the effectiveness of the penalty regime in deterring socially harmful actions: as a result, firms get away with investing in innovation more than it would be socially optimal. Unsurprisingly, this occurs for relatively high β and relatively low p, namely, when social harm is likely and enforcement is relatively ineffective.

7. Extensions

In this section we discuss four different extensions of the analysis developed so far, to see how its main predictions change in response to changes in assumptions. First, we show that the model can accommodate varying degrees of risk aversion in the social welfare function, and that a more risk-averse society will want to regulate innovative firms more stringently, especially by relying more on authorization than on penalties, although this comes at the cost of less investment in innovation. Second, we show that if the courts entrusted with the task of enforcing penalties on potential violators tend to incur frequently in type-II errors, society will again want to rely more heavily on authorization than on penalties, so as to reduce under-deterrence. Thirdly, we ask which would be the implications of allowing for patent races in the R&D investment decision by firms, and conclude that, to the extent that such races lead to more investment in R&D (just as an increased profitability of R&D activity), it would again lead regulators to choose less innovation-friendly policy regimes, since the concern of an inefficiently low investment in innovation would be reduced or absent. Finally, we consider how the model's predictions compare to those that one would obtain in the absence of costly investment in R&D: also in this case, being unconcerned with the need to sustain investment in innovation, policies will tend to be stricter, in that *laissez faire is* never optimal, and penalties for harmful actions are larger than in the presence of investment in R&D.

7.1. Social risk aversion

It is often argued that innovations may in some cases trigger catastrophic effects for society, which justify the adoption of a social welfare function with some degree of risk aversion. Indeed, the potentially large "downward risk" of innovations is often the main argument for extremely rigid rules.[14]

Our welfare function, being piecewise linear, can capture the concavity required by risk aversion by assuming its slope to be greater in the bad state than in the good one ($w^b > w^g$). This can be seen by rewriting our welfare function, currently defined as

$$W = \begin{cases} W^b = -w^b a & \text{in the bad state,} \\ W^g = w^g a & \text{in the good state,} \end{cases}$$

with $a \in (0,\bar{a}]$ in both states, as the piecewise linear function [15]

$$W = \begin{cases} w^b a & \text{for } a \in [-\bar{a}, 0] \text{ in the bad state,} \\ w^g a & \text{for } a \in [0, \bar{a}] \text{ in the good state.} \end{cases}$$

Risk-neutral preferences then correspond to $w^b = w^g$, and risk-averse preferences to $w^b > w^g$. Therefore, in our setting a higher degree of social risk aversion is captured by a larger w^b, for given w^g. It is easy to see that such a parameter shift leads to tougher, less innovation-friendly policies in all the cases analyzed so far. If fines are unbounded in the penalty regime, an increase in w^b, which makes actions in the bad state more welfare-decreasing, decreases the scope for *laissez faire* (Proposition 1). Moreover, with limited liability, a higher w^b shifts all the thresholds between regimes to the left, reducing the *laissez-faire* and penalty regions and correspondingly expanding the region of the authorization regime. Indeed, when

[13] A related issue is that of the distribution of enforcement costs, which is related to who bears the burden of proof. But if enforcement costs are funded by levying lump-sum taxes, it is irrelevant whether they are borne by the firm or by the government, since even in the latter case they are ultimately paid by the private sector. In either case, the firm will choose the same action and investment level, and the regulator will pick the same policy, since social welfare is the same.

[14] We thank one of the referees for making this point.

[15] Under this alternative formulation, the profit function becomes πa in the good state, when $a \in [0,\bar{a}]$, and $-\pi a$ in the bad state, when $a \in [-\bar{a}, 0]$.

$w^b > w^g - \pi$ – a condition which is satisfied if society is risk-averse – only authorization is chosen as the optimal regime, confirming the intuition that a risk-averse society should opt for very rigid rules.

7.2. Type-II errors and under-deterrence by courts

In contrast to what happens in authorization procedures, the judicial enforcement of laws is based on an *ex-post* investigation of actual behavior, as in our penalty regime. This investigation generally rests on the principle that the subjects under investigation are presumed innocent unless contrary evidence is uncovered. This principle implies that type II errors may occur frequently under the penalty regime, resulting in under-deterrence. This may be a substantial cost of the *ex-post* enforcement of penalties, compared to the *ex-ante* review of the evidence in authorization procedures (provided these are of the strict, rather than lenient, type).

To address this issue, assume that when an innovation is socially harmful the court with probability α_{II} misrepresents the true state of the world, committing a "type-II error" As a result, the firm choosing a harmful action is fined only with probability $p(1-\alpha_{II})$, so that in the bad state the firm chooses the action $\hat{a}$ that maximizes its profits, net of the expected fine:

$$\hat{a} = \arg\max_{a \in A} [\pi a - p(1-\alpha_{II}) f(a)].$$

The corresponding investment in innovative activity is

$$I_P^{\alpha} = \frac{(1-\beta)\pi\bar{a} + \beta[\pi\hat{a} - p(1-\alpha_{II})f(\hat{a})]}{c},$$

which is increasing in α_{II}, due to under-deterrence. Introducing a type-II error modifies the thresholds between the penalty and lenient authorization regimes in the limited liability case considered in Proposition 2. It is easy to show that the threshold for lenient authorization then becomes:

$$\beta_P(p, \alpha_{II}) = \frac{w^g - w^b - \pi}{w^g + w^b - p(1-\alpha_{II})(2w^b + \pi)},$$

which is decreasing in α_{II}. Hence, the region where penalties are chosen (the interval where $\beta < \beta_P(p, \alpha_{II})$) shrinks to the advantage of the authorization regime when the under-deterrence of the penalty regime stemming from type-II errors becomes a more serious concern.[16]

7.3. Patent races

Our model relies on a simple description of investment in innovation that is consistent with the approach by Arrow (1962), whereby a single firm can invest in R&D. In this setting, over- or under-investment may arise due to the misalignment of private and social incentives. In practice, often firms are seen to engage in competition to develop an innovation, i.e. patent races. Then the question arises as to whether our setting can be consistent with such patent races. The literature has considered two different paradigms of patent races, which also produce the private–public misalignment just mentioned, but introduce two further reasons for over-investment. In Dasgupta and Stiglitz (1980, Section 3) firms compete à la Bertrand to win the race, which may lead them to engage in excessive individual investments. In Lee and Wilde (1980), Loury (1979) and Dasgupta and Stiglitz (1980, Section 4), instead, uncertainty gives a chance of winning to all participants, inducing excessive participation. These additional reasons for over-investment are not present in our model. Hence, by combining our approach with a different modelling of patent races we may have more investment in innovation than that obtained in the present setting.

The implication of this higher level of investment in R&D for optimal public policy is that it will lean towards more rigid forms of intervention. The intuition is that a higher level of R&D activity makes the regulator less concerned with the goal of sustaining R&D activity and therefore re-orients policy towards a stricter control of the actions chosen by the firm. This can be seen, for instance, by considering the policy reaction to a higher propensity to invest as captured, for instance, by a parametric increase in π. First, it is immediate that investment in research is increasing in π in all regimes, as it is clear from inspecting Eq. (31) for I_r. Moreover, as shown below, all the thresholds between regimes are decreasing in π:

$$\frac{d\beta_{LF}}{d\pi} = -\frac{1}{w^g + w^b} < 0,$$

$$\frac{d\beta_P}{d\pi} = -\frac{\left(w^g + w^b\right)(1-p)}{\left[w^g + w^b - p(2w^b + \pi)\right]^2} < 0,$$

and

$$\mathrm{sgn}\left(\frac{d\beta_{LA}}{d\pi}\right) = \mathrm{sgn}\left(-1 + \frac{p\left(w^g - w^b - \pi\right)}{\sqrt[2]{\left(w^g + w^b\right)^2 + p^2\left(\pi^2 - 2\pi w^g + 2\pi w^b - 4w^g w^b\right)}}\right) \le 0$$

for any p.[17]

Hence, for given β, an increase in the level of R&D investment, induced for instance by patent races, tends to shift the optimal policy toward more rigid rules. Intuitively, this result rests on the idea that when private incentives to innovation are stronger, public policies should be less concerned with boosting innovative activity, and therefore become more interventionist.

7.4. The model without investment in research

It is worth comparing the results obtained so far with those that would arise in a setting where firms can implement the actions in A without undertaking costly investments in innovation. With this change in assumptions, the model reduces to a standard model of law enforcement, where policy is simply directed to contain the potential social harm stemming from the firm's choice among the actions in A. In this simpler version of the model, one must modify the time line by dropping stage 1, where firms choose their investment in R&D, in all the regimes analyzed in the previous section. As we will see, this change in assumptions implies a tougher policy stance against the social harm caused by the firms' actions, since the policy maker needs no longer worry of preserving the firm's incentives to invest in innovation. We consider the same regimes as in previous sections, and for each we compute the associated expected welfare, so as to rank them. In the *laissez-faire* regime, the expected welfare is simply that associated with action $\bar{a}$ by firms, that is,

$$E(W_{LF}) = (1-\beta)w^g\bar{a} - \beta w^b\bar{a}, \tag{23}$$

[16] For a more general treatment of type-I and type-II errors and the choice of the optimal level of accuracy and legal standard in an antitrust setting, see Immordino and Polo (2011).

[17] To show this, notice that the fraction in the parenthesis is weakly smaller than 1 if $p(w^g - w^b - \pi) \le \sqrt[2]{(w^g + w^b)^2 + p^2(\pi^2 - 2\pi w^g + 2\pi w^b - 4w^g w^b)}$, which is satisfied for $p = 1$. But since the left-hand side of this inequality is increasing in p, while the right-hand side is decreasing in p, this inequality is satisfied for any p. Hence, the fraction never exceeds 1.

whereas in the two authorization regimes, it turns out to be respectively

$$\mathrm{E}(W_{LA}) = (1-\beta)w^g\bar{a}-\beta(1-p)w^b\bar{a} \tag{24}$$

and

$$\mathrm{E}(W_{SA}) = (1-\beta)pw^g\bar{a} \tag{25}$$

Notice that, as in Section 4, lenient authorization is preferred to the strict authorization if the bad state is sufficiently unlikely, that is Eq. (24) exceeds Eq. (25) if $\beta \le \beta_{LA}' = w^g/(w^g+w^b)$.

In the penalty regime, expected welfare is $\mathrm{E}(W_P) = (1-\beta)w^g a_P^g - \beta(1-p)w^b a_P^b$ and the fine schedule $f(a)$ is set so as to implement the welfare-maximizing actions in each state. In the good state, clearly $a_P^g = \bar{a}$: welfare-enhancing actions are lawful. In the bad state, instead, when the maximum fine is unbounded, it is optimal to always implement the *status-quo* action $a_P^b = a_0$. Absent any indirect effect through innovative investment, allowing a more harmful action has only a negative impact on welfare: $\partial E(W_P)/\partial a_P^b = -\beta(1-p)w^b < 0$. In other words, average deterrence disappears and only marginal deterrence matters. The *status-quo* action is implemented through the fine schedule $f(a) \ge \frac{\pi\bar{a}}{p}$ for any $a > 0$. Therefore the expected welfare associated with the optimal policy is

$$\mathrm{E}(W_P) = (1-\beta)w^g\bar{a} \tag{26}$$

which corresponds to the first best and therefore exceeds $\mathrm{E}(W_{LF})$, $\mathrm{E}(W_{LA})$ and $\mathrm{E}(W_{SA})$.

When instead the maximum fine is capped at $\pi\bar{a}$ due to limited liability, the optimal action implemented in the bad state is $\bar{a}(1-p)$: setting the fine $\bar{f}$ at the highest feasible level $\pi\bar{a}$, the incentive compatibility constraint becomes

$$\pi\hat{a} = \pi\bar{a} - p(\pi\bar{a}) \Leftrightarrow \hat{a} = \bar{a}(1-p).$$

The associated welfare level is

$$\mathrm{E}(W_P) = (1-\beta)w^g\bar{a}-\beta(1-p)w^b\bar{a} = \mathrm{E}(W_{LA}). \tag{27}$$

Therefore, lenient authorization and penalty are now equivalent and both of them dominate *laissez faire*, as we can see by directly comparing Eqs. (23) and (27).

Summarizing:

Proposition 4. Optimal policy without innovative activity

When new actions do not require innovative activity and the maximum fine F is unbounded*, the penalty regime implements the first best through complete deterrence in the bad state. When instead a limited liability rule constraints the maximum fine, for $\beta \le \beta_{LA}'$ lenient authorization and penalty are equivalent and dominate both laissez faire and strict authorization, while for $\beta > \beta_{LA}'$ strict authorization is the optimal policy.*

Comparing the results obtained when innovative activity matters with the present environment, we can see that the optimal policy becomes tougher in two dimensions. First, there is no longer a region where the regulator is willing to abstain from any intervention, even when the probability of the bad state is very low: *laissez faire* is never optimal. Second, the penalties applying to harmful actions are larger, so as to achieve complete deterrence. Therefore, in the bad state the implemented action is invariably the *status quo* a_0, whereas in the presence of innovative investment it is optimal to leave firms some scope to choose more profitable and harmful actions.

8. Conclusion

In this paper we develop a model where firms can invest in costly research and then, if successful, undertake new types of production that, though profitable, may prove harmful to society. In these situations public policies have a direct effect on the new activities carried out by successful innovators and an indirect effect on their incentives to invest in innovation. Interventions should be designed so as to balance the prevention of social harm with the benefits from innovative activity. That this is an empirically relevant tradeoff is witnessed by evidence from the pharmaceutical industry, which shows that imposing tighter standards in the tests required for the introduction of new drugs has considerably slowed down the pace of innovation in this sector.

We suppose that regulators can intervene by picking one of four different regimes: (i) *laissez faire*, (ii) a *lenient authorization* regime where firms are allowed to exploit all innovations not proven to be harmful, (iii) a *strict authorization* regime where only innovations proven to be safe are allowed, and (iv) a *penalty* regime where firms can exploit their innovations commercially but are subject to fines if these are found to be socially harmful.

Our key result is that the regulatory regime should become increasingly stringent as the danger of social harm increases – a principle that appears consistent with the policies implemented in the pharmaceutical industry, as well with current proposals for the regulation of financial innovation. In particular, we show that when the danger of social harm is very low, *laissez faire* is optimal, while above a critical risk threshold the regulator should switch to a system of penalties, fine-tuning them to induce firms to choose less and less damaging actions as the probability of harm increases. This result applies if fines can be always set high enough to induce firms to choose the welfare-maximizing action: in this case penalties are the most effective form of public intervention, while authorization, due to its yes-or-no nature, is not equally effective in controlling the firm's actions. Underinvestment in research occurs in all cases, as private incentives are lower than social welfare.

The picture changes if fines are subject to a limited liability constraint, so that they cannot exceed the firms' realized profits. In this case, when social harm is sufficiently likely, even the largest fine that can be inflicted on a non-complying firm cannot induce it to choose the more moderate action required by welfare maximization, so that the penalty regime becomes less effective than one based on authorization. In this case, if one considers increasing values of the likelihood of social harm, the sequence of optimal policies is *laissez faire*, penalties, lenient authorization and eventually strict authorization. Moreover, whenever lenient authorization is optimal, so are penalties: in that parameter region, the two regimes are equivalent. In a portion of that region, firms will overinvest in research, as a result of the under-deterrence caused by the limited liability constraint.

Our general result suggests that public policies should be softer when innovation is an important source of welfare improvements. This is further appreciated comparing the results just outlined with the optimal policies when no costly investment in research is required for firms to engage in new and potentially harmful activities. In that case, the prescribed policies entail a stricter control of socially damaging actions: *laissez faire* is never optimal, and penalties are geared to achieve complete deterrence, thus preventing any socially harmful action.

We show that a similar prescription applies if society becomes more risk-averse, which may be justified when innovation may cause catastrophic consequences if it turns out to be harmful: in this case, the scope for *laissez faire* become smaller, and in the presence of limited liability, an authorization regime becomes preferable to a system of penalties. The same tendency to opt for very rigid rules can be justified if the courts that are called to enforce penalties are biased towards acquitting potential violators when the evidence is indecisive, so that a system of penalties generates under-deterrence.

Acknowledgments

We are indebted to Maurizio Cafagno, Francesco Denozza, Luigi Franzoni, Nuno Garoupa, Michele Grillo, Dilip Mookherjee, Roberto Pardolesi, Franck Portier, Patrick Rey, Jean Charles Rochet, Aldo Schiavello, Andrei Shleifer, Jean Tirole, Aldo Travi, Helder Vasconcelos, and Georg von Wangenheim for their helpful comments and discussions, and the editor and two anonymous referees for the very valuable suggestions. We are also grateful to seminar participants in Bologna, Padua, Toulouse, Turin, Bocconi, the 2005 CSEF-IGIER Symposium on Economics and Institutions (Capri), the first Italian Congress in Law and Economics (Siena), the 2006 EEA Congress (Vienna) and the 2007 ESEM Congress (Budapest). This is a substantially revised version of an earlier paper titled "Norm Flexibility and Innovative Activity". We acknowledge financial support from MIUR and Università Bocconi.

Appendix A

Proof of Lemma 1. When $p=1$, we have Eq. (6) = Eq. (8) for any β. When $p<1$, equating the two expressions, solving for β and choosing the only root in $[0,1]$ yields

$$\beta_{LA}(p)=\frac{1}{2p(w^g+w^b)}\times\left(w^g+w^b-\sqrt{(w^g+w^b)^2+p^2(\pi^2+2\pi w^b-2\pi w^g-4w^gw^b)}+p(2w^g-\pi)\right), \tag{28}$$

which is increasing in p for $p\in[0,1]$ and tends to 1 when $p\to1$. Evaluating the expected welfare in the two regimes for $\beta=0$ and $\beta=1$ we have $E(W_{LA}(0))>E(W_{SA}(0))$ and $E(W_{LA}(1))<E(W_{SA}(1))$. Indeed

$$E(W_{LA}(0))=\frac{\pi\bar{a}^2}{c}\left(w^g-\frac{\pi}{2}\right)>E(W_{SA}(0))=\frac{\pi\bar{a}^2}{c}p^2\left(w^g-\frac{\pi}{2}\right)$$

and

$$E(W_{LA}(1))=-\frac{\pi\bar{a}^2}{c}(1-p)^2\left(w^b+\frac{\pi}{2}\right)<E(W_{SA}(1))=0.$$

Since $\beta_{LA}(p)$ is the only value of β for which the two functions coincide, for $\beta<\beta_{LA}(p)$ the *LA* regime dominates the *SA* one, and viceversa. ■

Proof of Lemma 4. Suppose $w^g-w^b-\pi>0$.

The maximization program must satisfy the conditions in Eqs. (16) and (17). We derive the solution for each of the two intervals in the lemma.

(i) For $\beta\in[0,\beta_{LF})$, where $\beta_{LF}>0$ is the value of β such that $\partial E(W_P)/\partial\tilde{a}=0$ for $\hat{a}=\bar{a}$, $\underline{f}=0$ and $\lambda=0$, we have

$$\frac{\partial E(W_P)}{\partial\tilde{a}}=\bar{a}\left[(1-\beta)w^g-(1+\beta)w^b-\pi\right]\frac{\beta\pi}{c}\geqslant0, \tag{29}$$

because the term in square brackets is non-negative (being decreasing in β and zero for $\beta=\beta_{LF}$) and the term $\beta\pi/c$ is also non-negative for $\beta\in[0,\beta_{LF})$. Hence, the regulator chooses the highest feasible action, $\hat{a}=\bar{a}$. Since in the same interval $E(W)-cI_P>0$, it follows that $\partial E(W_P)/\partial\underline{f}<0$. Hence, $\underline{f}=0$. Moreover, the incentive compatibility constraint does not bind, so that $\lambda=0$. In fact, if it were $\lambda>0$, then $\bar{f}=F$ and λ should be zero to satisfy the complementary slackness condition, leading to a contradiction. Since $\lambda=0$, the high fine $\bar{f}$ can be any value satisfying the incentive compatibility constraint.

(ii) For $\beta>\beta_{LF}$, the expected welfare from the innovative activity $E(W)$ is reduced and we obtain an internal solution from the first-order condition with respect to $\tilde{a}$. To show this, notice that if the regulator were to choose $\hat{a}=\bar{a}$ also in this interval, one would have $\partial E(W_P)/\partial\tilde{a}<0$, because the term in square brackets in Eq. (29) would be negative. Hence in this interval the regulator will choose $\hat{a}<\bar{a}$. The internal solution for the first-order condition implies that $E(W)-cI_P>0$ for $\lambda=0$, so that $\partial E(W_P)/\partial\underline{f}<0$. Hence, $\underline{f}=0$. Finally, also in this interval the incentive compatibility constraint is not binding, by the same argument used for the previous interval. Then, substituting $\underline{f}=0$ and $\lambda=0$ in $\partial E(W_P)/\partial\tilde{a}=0$ we get the welfare-maximizing implemented action $\hat{a}$ in Eq. (19), which is decreasing in the likelihood of the bad state β. Finally, to identify the high fine $\bar{f}$, we substitute $\hat{a}$ from Eq. (19) into the incentive compatibility constraint holding with equality. This yields the lower bound $\bar{f}(\beta)$ stated in the lemma. We conclude the proof by noticing that Eq. (19), together with $\underline{f}=\lambda=0$, identifies a unique maximum because $\partial^2E(W_P)/\partial\tilde{a}^2<0$.

Notice that when $w^g-w^b-\pi\leqslant0$ we obtain $\beta_{LF}\leqslant0$ and only case (ii) matters. ■

Proof of Proposition 1. First of all, recall from the previous analysis that if $w^g-w^b-\pi>0$, the penalty regime includes *laissez-faire* (*LF*) for $\beta\leqslant\beta_{LF}$, that is, no fine is inflicted for any action $a\in[0,1]$; if instead $w^g-w^b-\pi\leqslant0$, positive penalties are inflicted for actions above $\hat{a}$ for any $\beta\in[0,1]$. This immediately establishes that the *LF* regime occurs only when $w^g-w^b-\pi>0$.

Next, in order to identify the optimal policy for $\beta\in[0,1]$ two preliminary observations are important:

(i) In equilibrium the firm never pays a fine under any policy rule. This is obviously true in the authorization regimes, where no fine can be charged. But it is true as well in the *LF* regime, where no fine is set, and in the *P* regime, since the optimal fine schedule requires to set $\underline{f}=0$ at the implemented action $\hat{a}$. Hence, fines do not affect the expression of the innovative investment and of expected welfare when these are evaluated at the optimal policy in the different regimes.

(ii) The policy rules and the timing of their application affect differently the choice of the actions across regimes. Under *LF* the firm adopts the same action $\bar{a}$ in the good and bad state, in the *P* regime it adopts a different action according to the state of nature while in the *LA* and *SA* regimes the random outcome of the preliminary evaluation further determines different contingencies in the good and bad state. For instance, under the *LA* regime, in the bad state the firm will be authorized with probability $1-p$, choosing the most profitable action $\bar{a}$, while it will be forced to the *status quo* action if evidence of welfare losses is obtained. Hence, we cannot assign a single action to the bad state, although we can define the expected action in the bad state as $(1-p)\bar{a}$. However, given the linearity of the welfare functions, the expected welfare in the bad state can be written equivalently as a function of the actions chosen when the preliminary evaluation fails to produce hard evidence ($\bar{a}$) and when it does (a_0), or in terms of the expected action: $E(W^b_{LA})=-[(1-p)w^b\bar{a}+pw^ba_0]=-w^b(1-p)\bar{a}$.

These two observations allow us to conclude that in any policy regime r (for $r=LF,P,LA,SA$), expected welfare evaluated at the optimal policies can be written, for given β and p, as a function of the chosen (expected) actions in the good (a^g_r) and bad (a^b_r) state, making the corresponding functions perfectly comparable. The same holds for the investment chosen by firms. Hence, we can write in general the expected welfare associated to the regime r as

$$\mathrm{E}(W_r) = I_r\left[(1-\beta)w^g a_r^g - \beta w^b a_r^b\right] - c\frac{I_r^2}{2} \tag{30}$$

and the investment chosen in regime r as

$$I_r = \frac{(1-\beta)\pi a_r^g + \beta a_r^b}{c}. \tag{31}$$

The actions chosen in the good and bad state, identified in the previous analysis, are:

$$a_{LF}^g = \bar{a} \quad \text{and} \quad a_{LF}^b = \bar{a},$$

$$a_P^g = \bar{a} \quad \text{and} \quad a_P^b = \bar{a}\frac{(1-\beta)}{\beta}\frac{\left(w^g - w^b - \pi\right)}{(2w^b + \pi)} = \bar{a}\frac{(1-\beta)}{\beta}\frac{(\beta_0)}{(1-\beta_0)},$$

$$a_{LA}^g = \bar{a} \quad \text{and} \quad a_{LA}^b = (1-p)\bar{a} + pa_0 = (1-p)\bar{a},$$

$$a_{SA}^g = p\bar{a} + (1-p)a_0 = p\bar{a} \quad \text{and} \quad a_P^b = a_0 = 0.$$

Notice that, since fines are unbounded in the P regime (that encompasses the LF policy), the agency can always implement the actions (including $\hat{a}$ in the bad state) that maximize the expected welfare in Eq. (30) with no constraint. In the authorization regimes, instead, the actions simply derive from the application of the corresponding authorization rules or reflect profit maximization in the contingencies in which the firm obtains green light to its request. In other words, the actions taken in the authorization regimes are not the result of a complete welfare maximization program. Hence, the (LF and) P regime by definition (weakly) dominates the LA and SA regimes. The latter may give the same expected welfare as the P regime for particular values of β and p as long as the actions taken in each contingency by the firm are the same as those in the P regime.

Hence, to check whether the expected welfare under the authorization regimes may be equal to the maximum welfare, which corresponds to that obtained in the P regime, we compare the actions chosen in the good and bad states. Consider first the LA and P cases. In the good state the two regimes lead to the same action $\bar{a}$ while in the bad state they generally lead to different (expected) actions. To investigate if there are some combination of parameters (p,β) that induce the choice of the same (expected) actions in the bad state in the two regimes, we equate the action chosen under the P regime (a_P^b) to a_{LA}^b:

$$\bar{a}\frac{(1-\beta)}{\beta}\frac{\left(w^g - w^b - \pi\right)}{(2w^b + \pi)} = (1-p)\bar{a}$$

and solve for β:

$$\beta_P(p) = \frac{w^g - w^b - \pi}{w^g - w^b - p(2w^b + \pi)} > \beta_{LF}. \tag{32}$$

At $(\beta_P(p),p)$ the two regimes are welfare-equivalent, while for any other combination of parameters (LF and) P strictly dominates LA, since the actions chosen under LA do not correspond to the (welfare maximizing) actions implemented under P. Notice that $\beta_P(p)$ is increasing in p and ranges from β_{LF} when $p=0$ to 1 when $p=1$. Consider next the comparison of P and SA, by comparing the corresponding actions chosen. For $\beta=1$ only the bad state action matters under both P and SA and the *status-quo* action $a_0=0$ is chosen both in P and SA regimes, making them welfare-equivalent. For $\beta<1$, instead, the actions chosen in the good and bad states differ in the two regimes. Since in both regimes expected welfare is obtained from the same expression in Eq. (30) computed at the chosen actions, P, being the result of maximizing Eq. (30), dominates SA. Summing up, we conclude that for any $\beta\in[0,1]$ (LF and) P weakly dominate LA and SA, that is $\mathrm{E}(W_P)\geqslant\max\{\mathrm{E}(W_{LA}),\mathrm{E}(W_{SA})\}$, with the strict inequality sign holding for $\beta\neq\{\beta_P(p),1\}$. ■

Proof of Proposition 2. Let us use for simplicity the notation $\beta_P(p)=\beta_P$. We first establish that, when $p<1$, the following ranking among thresholds holds:

$$\beta_{LF} \leqslant \beta_P < \beta_{LA}.$$

The corresponding expressions are given by Eqs. (18), (32) and (28) respectively. The former inequality ($\beta_{LF}\leqslant\beta_P$) derives from direct comparison of the corresponding expressions, the strict inequality holding for $p>0$. For the latter inequality ($\beta_P<\beta_{LA}$) direct comparison is difficult, but we can use the following argument. First, we know from its definition that at β_P the expected welfare in the P and LA regimes are equal, i.e. $\mathrm{E}(W_{LA}(\beta_P))=\mathrm{E}(W_P(\beta_P))$. Second, β_{LA} is defined as the value that equates the expected welfare in the SA and LA regimes, that is $\mathrm{E}(W_{LA}(\beta_{LA}))=\mathrm{E}(W_{SA}(\beta_{LA}))$. Moreover, we know that $\beta_{LA}<1$ (Lemma 1) and $\beta_P<1$ (Proposition 1) when $p<1$. Finally, we have established in Proposition 1 that P dominates SA for $\beta<1$. Putting all these facts together, we have that $\mathrm{E}(W_{LA}(\beta_P))=\mathrm{E}(W_P(\beta_P))>\mathrm{E}(W_{SA}(\beta_P))$, that is at β_P LA dominates SA. Hence it must be that $\beta_P<\beta_{LA}$. We can now turn to the optimal fine in the P regime under limited liability.

In the optimal fine schedule the incentive compatibility constraint requires that the maximum fine be set according to the inequality in Eq. (20). Evaluating it at $\beta=\beta_P$ we obtain

$$F \geqslant \bar{f} \geqslant \frac{\pi\bar{a}}{p}\frac{\beta_P\left(w^g + w^b\right) - \left(w^g - w^b - \pi\right)}{\beta_P(2w^b + \pi)} = \pi\bar{a}.$$

Hence, at $\beta=\beta_P$ the limited liability constraint $F=\pi\bar{a}$ starts binding and the maximum fine cannot be increased further, nor the implemented action $\hat{a}$ be decreased. Consequently, for any $\beta>\beta_P$ in the P regime the implemented action is constant and corresponds to $\hat{a}(\beta) = \hat{a}(\beta_P) = \bar{a}(1-p)$, the same (expected) action chosen under the LA regime, as noted in the proof of Proposition 1. Since for any $\beta\geqslant\beta_P$ the actions chosen in the P and LA regimes are the same in the good and in the bad state for any $\beta\geqslant\beta_P$, the two regimes are equivalent and feature the same expected welfare: $\mathrm{E}(W_P(\beta))=\mathrm{E}(W_{LA}(\beta))$. Moreover, since for $p<1$ we have $\beta_P<\beta_{LA}<1$ as shown above and in Lemma 1, LA dominates SA for $\beta\in(\beta_P,\beta_{LA}]$, while the opposite holds true for $\beta\in(\beta_{LA},1]$.For $p=1$ we have $\beta_P=1$ and the last two intervals disappear, so that authorization regimes are never chosen.

Finally, notice that for $p<1$, if $w^g-w^b-\pi>0$, then $0<\beta_{LF}\leqslant\beta_P<\beta_{LA}<1$ and all the four regimes are chosen for $\beta\in[0,1]$; instead, if $w^g-w^b-\pi\leqslant0$ we have $\beta_{LF}\leqslant\beta_P<0<\beta_{LA}<1$ and only lenient and strict authorization are selected for $\beta\in[0,1]$. ■

Proof of Proposition 3. Recall that innovative investment in the different policy regimes is given by the same general expression in Eq. (31). Once evaluated at the corresponding actions, described in the proof of Proposition 1, we obtain

$$I_{FB} = \frac{\pi\bar{a}}{c}\frac{w^g}{\pi}(1-\beta); \quad I_{LF} = \frac{\pi\bar{a}}{c}; \quad I_P = \frac{\pi\bar{a}}{c}\frac{1-\beta}{1-\beta_{LF}};$$

$$I_{LA} = \frac{\pi\bar{a}}{c}(1-p\beta); \quad I_{SA} = \frac{\pi\bar{a}}{c}(1-\beta)p.$$

Moreover, Propositions 1 and 2 identify the parameter regions where the various regimes are optimal.

Let us consider first the case in which the maximum fine F is unbounded. Then, it is optimal to adopt LF for $\beta\in[0,\beta_{LF}]$ and P for

$\beta \in (\beta_{LF}, 1]$. Consider first the interval $\beta \in [0, \beta_{LF}]$. The first-best investment is larger than *laissez-faire* investment at $\beta = 0$: $I_{FB}(0) = \frac{w^g \bar{a}}{c} > I_{LF} = \frac{\pi \bar{a}}{c}$. Moreover, the first-best investment $I_{FB}(\beta)$ is decreasing and linear while I_{LF} is constant in the interval $[0, \beta_{LF}]$. Then, in order to check if the first-best investment is larger than the *laissez-faire* one over this interval, it is sufficient to check that the former is larger at boundary of the interval, that is at $\beta = \beta_{LF}$. Substituting and rearranging we obtain that $I_{FB}(\beta_{LF}) > I_{LF}$ if $2w^g > \pi$, which clearly holds. Hence, for $\beta \in [0, \beta_{LF}]$ the optimal *LF* policy yields underinvestment in innovation.

Moving to the interval $\beta \in (\beta_{LF}, 1]$, it is optimal to adopt regime *P*. Evaluating the corresponding investment at the boundary of this interval at $\beta = \beta_{LF}$, we obtain

$$I_P(\beta_{LF}) = \frac{\pi \bar{a}}{c} = I_{LF},$$

that is, investment varies continuously at the boundary between $[0, \beta_{LF}]$ and $(\beta_{LF}, 1]$. Hence, $I_P(\beta_{LF}) < I_{FB}(\beta_{LF})$. Moreover, looking at the corresponding expressions we can see that the I_{FB} and I_P functions are decreasing and linear and coincide at $\beta = 1$, since $I_P(1) = I_{FB}(1) = 0$. Hence, the optimal second-best policy entails underinvestment for any $\beta \in [0, 1)$.

Consider now the case of limited liability discussed in Proposition 2, which entails a richer set of optimal policies. For $p < 1$, indeed, welfare maximization requires to select *LF* for $\beta \in [0, \beta_{LF}]$, *P* for $\beta \in (\beta_{LF}, \beta_P]$, equivalently *P* or *LA* for $\beta \in (\beta_P, \beta_{LA}]$ and *SA* for $\beta \in (\beta_{LA}, 1]$. The previous proof that *LF* and *P* yield underinvestment extends to the limited liability case, but apply only to the interval $\beta \in [0, \beta_P]$, since the *P* regime is no longer optimal for $\beta > \beta_P$.

Hence, to complete the analysis we must consider investment in the *LA* and *SA* regimes for $\beta > \beta_P$. Since at $\beta = \beta_P$ the actions in the *P* and *LA* regimes are the same , the corresponding investment will be the same, and the investment associated to the optimal regimes is continuous at the boundary of this interval, i.e. $I_{LA}(\beta_P) = I_P(\beta_P)$. Therefore, by the previous discussion we have $I_{LA}(\beta_P) < I_{FB}(\beta_P)$.

Investment under *LA* is decreasing and linear in β and exceeds the first-best investment at $\beta = 1$: $I_{LA}(1) = \frac{\pi \bar{a}}{c}(1-p) > I_{FB}(1) = 0$. Therefore, the two functions intersect once in the interval $[\beta_P, 1]$. Equating $I_{LA}(\beta)$ and $I_{FB}(\beta)$ and rearranging we obtain that they coincide for β defined by the function

$$\beta_U = \frac{w^g - \pi}{w^g - p\pi},$$

i.e. $I_{LA}(\beta_U) = I_{FB}(\beta_U)$ and for $\beta > \beta_U$ *LA* implies overinvestment. Notice that β_U is increasing in *p*. Moreover for any *p* we have $\beta_U > \beta_P(p)$ if $2w^g - \pi > 0$, that clearly holds. Hence, the threshold β_U is in the region where *LA* replaces (or is equivalent to) *P*. Therefore we need to check whether *LA* or *SA* is chosen at $\beta = \beta_U$, that is, whether $EW_{LA}(\beta_U) \gtreqless EW_{SA}(\beta_U)$. Since $I_{SA} < I_{FB}$ for any $\beta < 1$, if at $\beta = \beta_U$ the optimal policy is *SA*, underinvestment occurs, while in the opposite case we would have overinvestment as long as *LA* is chosen, i.e. for $\beta \in (\beta_U, \beta_{LA})$. Substituting β_U in the two expressions for expected welfare we obtain, after rearranging, that the condition $EW_{LA}(\beta_U) \gtreqless EW_{SA}(\beta_U)$ is equivalent to the condition $p \gtreqless p_U$, where p_U is defined by

$$p_U = \sqrt{\frac{w^g \left(w^g \pi - 2 w^g w^b + 2 w^b \pi \right)}{\pi^2 (2 w^g - \pi)}}.$$

Summing up, for $p < p_U$ there is an interval $\beta \in (\beta_U, \beta_{LA})$ in which *LA* is the optimal policy regime and in which it leads to overinvestment. Instead for $p \geq p_U$ in the whole interval $\beta \in (\beta_P, \beta_{LA})$ where *LA* is optimal, there is underinvestment, a pattern that continues also for $\beta \in [\beta_{LA}, 1]$ once the optimal policy switches to *SA*. ■

References

Acemoglu, D., Verdier, T., 2000. The choice between market failures and corruption. The American Economic Review 90, 194–211.

Armstrong, M., Sappington, D.E.M., 2007. Recent developments in the theory of regulation. In: Armstrong, M., Porter, R.H. (Eds.), Handbook of Industrial Organization, Volume 3. Elsevier, Amsterdam, pp. 1557–1700.

Arrow, K.J., 1962. Economic welfare and the allocation of resources for invention. In: Nelson, R.R. (Ed.), The Rate and Direction of Inventive Activity: Economic and Social Factors. Princeton University Press, Princeton, pp. 609–626.

Banerjee, A.V., 1997. A theory of misgovernance. Quarterly Journal of Economics 112, 1289–1332.

Becker, G.S., 1968. Crime and punishment: an economic approach. Journal of Political Economy 76, 169–217.

Becker, G.S., Stigler, G.J., 1974. Law enforcement, malfeasance and compensation of enforcers. Journal of Legal Studies 17, 65–81.

Biello, D., 2008. Government Fails to Assess Potential Dangers of Nanotechnology. Scientific American, December 18.

Blankfein, L., 2009. Do Not Destroy the Essential Catalyst of Risk. Financial Times, 8 February, 7.

Dasgupta, P., Stiglitz, J., 1980. Uncertainty, industrial structure, and the speed of R&D. The Bell Journal of Economics 11, 1–28.

Financial Times, 2010a. BIS Wants OTC Derivatives Rated like Drugs. 16 June.

Financial Times, 2010b. Pharmaceuticals: Perils for Pill Pushers. 21 September.

Glaeser, E.L., Shleifer, A., 2003. The rise of the regulatory state. Journal of Economic Literature 41, 401–425.

Grabowski, H.G., Vernon, J., Thomas, L.G., 1978. Estimating the effects of regulation on innovation: an international comparative analysis of the pharmaceutical industry. Journal of Law and Economics 21, 133–163.

Hayek, F., 1935. Collectivist economic planning: critical studies on the possibilities of socialism. Routledge, London.

Hayek, F., 1940. Socialist calculation: the competitive 'solution'. Economica 7, 125–149.

Immordino, G., Pagano, M., 2010. Legal standards, enforcement and corruption. Journal of the European Economic Association 8, 1104–1132.

Immordino, G., and Polo, M., 2011. Antitrust Policy, Legal Standards and Innovative Activity. Unpublished.

Krueger, A., 1974. The political economy of the rent-seeking society. The American Economic Review 64, 291–303.

Laffont, J.J., Tirole, J., 1993. A theory of incentives in procurement and regulation. MIT Press, Cambridge.

Lee, T., Wilde, L.L., 1980. Market structure and innovation: a reformulation. Quarterly Journal of Economics 94, 429–436.

Loury, G., 1979. Market structure and innovation. Quarterly Journal of Economics 93, 395–410.

Mookherjee, D., Png, I.P.L., 1994. Marginal deterrence in the enforcement of law. Journal of Political Economy 102, 1039–1066.

Philipson, T.J., Sun, E.C., 2008. Is the food and drug administration safe and effective? The Journal of Economic Perspectives 22, 85–102.

Philipson, T.J., Sun, E.C., Goldman, D., 2010. The effects of product liability exemption in the presence of the FDA. In: Kessler, D.P. (Ed.), Regulation vs. Litigation: Perspectives from Economics and Law. University of Chicago Press, Chicago, pp. 137–163.

Polinsky, A.M., Shavell, S.M., 2000. The economic theory of public enforcement of law. Journal of Economic Literature 38, 45–76.

Rose-Ackermann, S., 1978. Corruption: A Study in Political Economy. Academic Press, New York.

Schwartzstein, J., Shleifer, A., 2009. Litigation and regulation. NBER Working Paper No. 14752.

Scientific American, 2010. Big Need for a Little Testing. January, 23.

Segal, I., Whinston, M., 2007. Antitrust in innovative industries. The American Economic Review 97, 1703–1730.

Shavell, S., 1980. Strict liability versus negligence. The Journal of Legal Studies 9, 1–25.

Shavell, S., 1984. Liability for harm versus regulation of safety. Journal of Legal Studies 13, 357–374.

Shavell, S., 2007. Do excessive legal standards discourage desirable activity? Economics Letters 95, 394–397.

Stigler, G.J., 1970. The optimum enforcement of laws. Journal of Political Economy 78, 526–536.