Selective Incentives*

Giuseppe Dari-Mattiacci[†] and Alex Raskolnikov[‡]

[Preliminary — Do not distribute]

Abstract

The fundamental premise of the economic analysis of law enforcement is that rational agents respond to expected sanctions. Therefore greater expected sanctions lead to greater compliance. The foundational conclusion of the literature on positive and negative incentives is that carrots and sticks have identical incentive effects. Therefore the choice between these two instruments depends on second-order considerations. We extend the canonical model of law enforcement in three ways. We allow the agents to not only comply or violate the rule, but to opt out of the regulatory regime altogether. We introduce rewards into the enforcement scheme. And we model a population of agents with varying private benefits or perceptions of the principal's enforcement accuracy. We show that when the participation and the compliance decisions are considered jointly, when carrots exist alongside sticks, and when agents have varying benefits or accuracy perceptions, neither of the two fundamental insights just described holds. Greater expected sanctions do not necessarily lead to more compliance, and carrots and sticks are not substitutes in their incentive effects. We also show that adding taxes and subsidies to the regulatory toolkit does not expand the set of achievable outcomes. We conclude that regulators should abandon the standard view about the effects of expected sanctions and the interchangeability of carrots and sticks when regulated parties derive low private benefits from participating in the regulated activity and when they believe that the regulator is highly accurate in separating compliers from violators.

Keywords: deterrence, compliance, participation, accuracy, sanctions, carrots, sticks, incentives.

JEL codes: K10, K20, K42.

1 Introduction

Gary Becker's seminal article on the economic analysis of crime (Becker 1968) was not really about crime. Or at least not only about it. In the article's first paragraph Becker emphasized that he would consider not only murder, rape and similar felonies, but also discrimination, collusion, traffic safety, and "thousands of other activities." One page later he added tax evasion and white-collar offenses to the list. And by the time he turned to applications of his model, he quipped

*We received valuable feedback from Scott Baker, Ryan Bubb, Gerrit De Geest, Ezra Friedman, and workshop participants at Columbia University, University of Chicago, NYU, and the annual meetings of ComplianceNet and the American Law and Economics Association. All mistakes are ours.

† Alfred W. Bressler Professor of Law, Columbia University School of Law.

‡ Wilbur H. Friedman Professor of Tax Law, Columbia University School of Law.

that "crime is apparently not so different analytically from any other activity that produces external harm[,] and when crimes are punishable by fines, the analytical differences virtually vanish" (Becker 1968:201). Becker's model had a broad reach indeed.

The theory of optimal deterrence that grew out of Becker's foundational work embraced his ambition. That theory is also known as the theory of public enforcement of law (Polinsky and Shavell 2007)—a name that leaves no doubt that the theory's goal is to analyze <u>all</u> government regulation. Yet the theory's workhorse model serves as an adequate representation of only a small number of regulatory settings. The model's key limitation is that it incorporates one choice facing most regulated parties while largely ignoring another one.

The individual's decision at the center of the standard deterrence model is whether to comply with the legal rule or to violate it. Economic analysis of law enforcement has much to say about how the government can shape this decision. But compliance and violation are not the only options available to individuals facing most regulatory regimes. The third option is to avoid participating in the regime altogether. The choice of whether or not to participate in a regulated activity has received much less attention from law-and-economics scholars. The analysis of how individuals optimize along the participation and the compliance margins simultaneously is particularly thin.

Yet Becker's own illustrations reveal the importance of addressing both choices. Consider his first example—discrimination, say housing discrimination to be concrete. The obvious goal of anti-discrimination laws is to induce compliance with these laws by landlords. But discrimination is difficult to define and to prove, so enforcement is never perfectly accurate. Facing a possible mistaken imposition of liability, some landlords may decide to exit rental real estate business altogether, or may not enter it in the first place. Importantly, some these landlords would have run a perfectly non-discriminatory, socially beneficial business. Their decision to stay away is a social cost.

The same analysis applies to collusion—Becker's second example. It, too, is difficult to define and to prove. And here, too, a possible mistaken imposition of antitrust liability may induce some firms to exit concentrated industries where a collusion prosecution is more likely. Turning to traffic safety, speed limits are sometimes uncertain and are enforced with error (Brabender 2004). Fearing mistakenly imposed sanctions, some would-be drivers may switch to public transportation or just stay home. More similar examples are easily available, and they reveal two challenges faced by regulators in the vast majority of cases. The first challenge is to ensure compliance with legal rules by those who participate in a regulated activity. The second challenge is to induce participation in that activity (or to prevent current participants from opting out).[1] Ignoring the second challenge necessarily limits the usefulness of the model.

The voluminous literature on the use of positive and negative incentives suffers from a similar failure to incorporate an important feature of real-world incentive schemes. The economic analysis of carrots and sticks has established long ago that the two instruments produce identical incentive effects (Ben Shahar and Bradford 2012). In fact, carrots and sticks are so similar that some scholars argue that the two are actually one and the same (Gordon 19XX). This does not mean, of course, that real-world positive and negative incentives are fully interchangeable. Carrots are expensive; sticks less so. Carrots induce participation; sticks discourage it (Wittman 1984). Carrots are given

[1] We recognize that in some settings participation is not optional. As long as one belongs to the society, one is required to comply with its criminal laws. Even here, however, non-participation is sometimes an option. For instance, if one is concerned that accidentally taking someone else's umbrella on the way out of a restaurant may lead to a mistaken conviction for theft, one may stop eating out altogether.

to all eligible claimants; sticks may never apply because individuals may take action to avoid them. Carrots work better to incentivize complex behaviors; sticks work well in simple settings (De Geest and Dari-Mattiacci 2013). These are just some of the known differences. But importantly, all these difference are predicated on the basic, fundamental insight that in terms of their incentives, carrots and sticks are identical.

In this article, we show that the fundamental results just described do not hold as a general matter. Expected sanctions do not always improve compliance, and carrots and sticks do not always have identical incentive effects. Our conclusions come from extending the basic model of deterrence to include rewards and the choice to abstain from participation in the regulatory regime. Equivalently, our results come from extending the basic model of carrots and sticks to include imperfect enforcement in response to inappropriate claiming of carrots. In each case, agent heterogeneity and enforcement errors are key to our analysis.

To see why higher expected sanctions do not always improve compliance, consider a simple example. A principal enforces a regulation by using a penalty (or "stick") equal to $s = 1$ for violators. When enforcing the regulation, the principal is not perfectly accurate; she correctly assesses violations and compliance with probability $q = \frac{3}{4}$. Thus, there is a twenty five percent chance (probability $1 - q = \frac{1}{4}$) that the principal concludes that a complying agent violated the rule or that a violating agent complied. Compliance requires agents to exert effort at a cost, e,[2] that varies across agents. Assume that e is uniformly distributed between 0 and 1.

If the agent complies, he bears the cost of compliance plus the expected cost of mistaken penalties, or $e + (1 - q)s = e + \frac{1}{4}$. If the agent violates, he bears the expected cost of correctly imposed penalties, or $qs = \frac{3}{4}$. The agent will comply if the cost of compliance is less than the cost of violation, that is, if $e < \frac{1}{2}$. Given our assumption about the distribution of e, half of all agents comply.

What if the principal increases the penalty from $s = 1$ to $s = 2$? The costs of compliance and violation become $e + \frac{1}{2}$ and $\frac{3}{2}$, respectively. Given this higher penalty, an agent complies if $e < 1$, meaning that now all agents comply.[3]

That compliance goes up with the magnitude of the penalty in a rational-choice model seems obvious, and is probably the most fundamental tenet of law-enforcement theory. However, this obvious and fundamental conclusion turns out to hold only if the sole choice faced by the agents is whether to comply or to violate. If we allow agents to also choose whether to participate, the monotonic relationship between penalties and compliance breaks down.

To see why, let us now consider two subgroups in the population of agents: high-benefit agents, who derive a benefit $b_H = \frac{3}{2}$ from participation, and low-benefit agents who derive a benefit $b_L = \frac{1}{2}$ from participation. (In both subgroups, e is uniformly distributed between 0 and 1.) If the agent does not participate in the regulated activity, his payoff is equal to 0. By subtracting the cost of compliance or violation from the benefit of participation for each subgroup of agents we obtain the payoffs reported in Table 1.[4]

Among the three available options, each agent chooses the one associated with the highest payoff. All high-benefit agents participate, because the payoff of both compliers and violators is

[2] The effort e can be interpreted as a foregone benefit of violation. In the literature on law enforcement, the benefit of violating the rule is usually denoted as g and varies among individuals, as e varies in our framework. The two narratives are interchangeable.

[3] That is, compliance yields lower costs than violation if $e + \frac{1}{2} < \frac{3}{2}$.

[4] For instance, the payoff of a complying high-benefit agent is equal to $b^H - [e + (1 - q)s] = \frac{3}{2} - \left(e + \frac{1}{4}\right) = \frac{5}{4} - e$. Other payoffs are calculated in an analogous way.

	Compliance	Violation	No participation
High-benefit agents	$\frac{5}{4} - e$	$\frac{3}{4}$	0
Low-benefit agents	$\frac{1}{4} - e$	$-\frac{1}{4}$	0

Table 1: Payoffs with $s = 1$.

greater than 0. Thus high-benefit agents choose—as before—between compliance and violation, and comply if $e < \frac{1}{2}$. In contrast, a low-benefit agent earns a negative payoff if he participates and violates the rule. Hence only those low-benefit agents who would comply if they participated may decide to participate. A low-benefit agent's payoff from compliance is positive if $e < \frac{1}{4}$. Therefore, only $\frac{1}{4}$ of the low-benefit agents participate and comply, while $\frac{3}{4}$ of them—those with $e \geq \frac{1}{4}$—do not participate (Figure 1a).

What if the principal increases the penalty from $s = 1$ to $s = 2$? The parties' payoffs will change, and the new payoffs are summarized in Table 2. As before, all high-benefit agents choose to participate, but now all of them comply—a higher penalty induces more compliance among these agents. This change increases the number of complying agents relative to the initial, low-penalty setting. The response of low-benefit agents is very different. Facing a higher penalty, all low-benefit agents choose not to participate. This change reduces the number of complying agents compared to the low-penalty setting.

	Compliance	Violation	No participation
High-benefit agents	$1 - e$	0	0
Low-benefit agents	$-e$	-1	0

Table 2: Payoffs with $s = 2$.

The intuition is that, an increase in the penalty reduces the payoff from violation more than it reduces the payoff from compliance (compare Table 1 with Table 2). This change induces more high-benefit agents to comply. Low-benefit agents, however, do not compare compliance with violation-–they compare compliance with non-participation. The payoff from non-participation is fixed at 0 no matter what the penalty is, but enforcement errors reduce the payoff from compliance when the penalty increases. Given this lower payoff fewer agents choose to participate and comply.

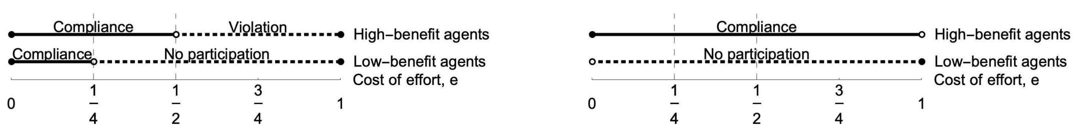

(a) Low penalty (b) High penalty

Figure 1: Agents' behavior

More generally, an increase in the penalty affects the compliance and the participation decisions in opposite ways. Depending on the relative proportions of high- and low-benefit agents in the population, increasing the penalty may lead to an overall reduction, rather than an increase, in the number of complying agents. In the example, if more than $\frac{2}{3}$ of the agents are low-benefit types, the decline in the number of compliers in the (relatively large) low-benefit group more than offsets the increase in the number of compliers in the (relatively small) high-benefit group. In fact, the overall reduction in compliers may exceed the reduction in violators in the high-benefit group. If that happens, a familiar policy of increasing penalties would result in an unfamiliar outcome of reduced compliance.[5]

Our second result modifies the view that carrots and sticks have identical incentive effects. The standard model of law enforcement generally focuses on two instruments-–the certainty and the severity of sanctions. A large literature on the use of positive and negative incentives–carrots and sticks–mostly ignores the problem of non-compliance, such as the improper claiming of carrots. In reality, however, carrots often exist alongside sticks, and regulators can (and do) use both tools-–imperfectly—to influence individual's decisions.

Returning to the housing discrimination example, the government has long-offered tax incentives to induce landlords to invest in low-income housing (Congressional Research Service 2019). Congress recently enacted another tax incentive program-–the Opportunity Zone tax credit—that may also subsidize construction of rental housing in economically disadvantaged areas (Tax Policy Center 2018). Needless to say, Congress enacted these incentive programs not to subsidize discriminatory low-income housing. But given the challenge of recognizing and prosecuting housing discrimination, both anti-discrimination laws and tax incentives affect decisions of would-be landlords—those who plan to discriminate and those who do not. More generally, carrots-–and not only sticks-–influence individuals' choices to participate and to comply with the law.

We show that a regulator's use of carrots leads to ambiguous results, just like a regulator's use of sticks does. More generous carrots may lead not only to more participation, but to more violations as well. Some would-be violators who would otherwise stay out of the regulated activity may decide to participate and capture the reward if the regulator mistakenly treats them as compliers. Once it is understood that sticks may result in less compliance and that carrots may result in more violations, it becomes clear that the two instruments are not necessarily substitutes, contrary to what is generally assumed.[6] As our general model will demonstrate, carrots and sticks are substitutes in the agents' choice between compliance and violation but are complements in the agents'

[5]Denote as h the fraction of high-benefit agents. With the original penalty of $s = 1$, total compliance is $\frac{1}{2}h + \frac{1}{4}(1-h) = \frac{1+h}{4}$. With the increased penalty of $s = 2$, total compliance is equal to h, as only the high-benefit agents comply and all of them do so. Therefore, increasing the penalty from 1 to 2 reduces compliance if $h < \frac{1+h}{4}$ that is, if $h < \frac{1}{3}$ and less than $\frac{1}{3}$ of the agents are high-benefit while more than $\frac{2}{3}$ are low-benefit types.

[6]Gordon (___) argued that carrots and sticks are indistinguishable from each other because a carrot can always be rewritten as a stick. A lesser version of the carrots-and-sticks equivalence contends that they are (perfect) substitutes (see, for a recent contribution, Ben-Shahar and Bradford 2012). A few contributions stress the differences between carrots and sticks while remaining faithful to the idea that policymakers use either carrots or sticks (Polinsky 1979, Wittman 1984, De Geest and Dari-Mattiacci 2013). Even authors favoring the joint use of carrots and sticks study the effect of the resulting incentives on a single margin (Andreoni 2003; Ben Shahar and Bradford 2012; Gilpatric 2009; Moldovanu 2012). In contrast, we study the effect of combining rewards and penalties to affect two separate decisions: whether or not to participate and, conditional on participation, whether or not to comply. Dari-Mattiacci (2009) investigates how agents make decisions along both of these margins but he studies the use of a single incentive. Kleven and Kopczuk (2011) were first to analyze government benefits (carrots) jointly with variations in complexity, which they modeled as costly increases in enforcement accuracy. They did not include sticks or heterogeneous perceptions of error rates as we do. Moreover, the cost of greater accuracy born by agents played a key role in their model while it is absent from ours. Shavell (2017) suggests that subsidies may be appropriate to account for positive externalities of activities regulated by strict liability or negligence regimes. He does not consider the effects of inaccurate enforcement that are key to our model.

participation decision.While sticks generally deter participation and carrots encourage it,[7] their overall effect on compliance and violations depends on the relative proportions of high-benefit agents—who make a "compliance choice"—and low-benefit agents—who make a "participation choice"—in the population.

Finally, we show that the effects of carrots, sticks, and the probability of detection that we identify cannot be changed by introducing taxes and subsidies. A tax (or subsidy), we show, amounts to simultaneously increasing the stick and reducing the carrot by a fixed amount, without expanding the set of outcomes that the regulator may achieve. Adding more instruments to the policy mix does not make a difference in terms of achievable outcomes. To control two margins—the participation and compliance decisions—it does not help to use more than two instruments—the carrot and the stick.

Our results relate to a large literature on law enforcement, but we build on a very small number of key contributions considering law enforcement in the presence of errors and optional participation.[8] Enforcement errors play an important role in our model. Without errors, compliers would be insensitive to changes in the stick and violators would be insensitive to changes in the carrot. Thus, most of our results depend on the occurrence of error with a positive probability, which, we believe, is a realistic depiction of real-life enforcement systems. Png (1986) was the first to introduce errors in a model of law enforcement, and Kaplow (1994) was first to formally study the interaction between accuracy and sanctions in a legal system. Participation decision is also crucial for our analysis. Png (1986) identified a negative effect of errors on participation and argued that carrots can be used to complement sticks to counter that negative effect.[9] Kaplow (2011) introduced chilling in the model of imperfectly accurate law enforcement. He studied how the possibility of a mistaken imposition of sanctions affects the optimal burden of proof. While these contributions provide important foundations for our inquiry, none of them considers heterogeneity in the agents' participation decision, which is crucial to our analysis.

More generally, the literature on the economics of law enforcement has considered the participation decision in order to determine whether individuals have incentives to engage in the optimal level of activity in threshold-based (also known as negligence or fault-based) and strict liability regimes (Polinsky and Shavell 2007: 425). This literature addresses the question of deterrence—that is the choice between compliance and violation—conditional on participation. Relaxing this assumption drives our results.

The plan for the remaining of this paper is as follows. In Section 2 we will introduce our general model of regulatory enforcement with carrots and sticks, of which simpler models with only the carrot or the stick are special cases. We will then demonstrate our results in a simplified version of the model, which captures all of the main insights, and resort to the Appendix for the general proofs. In Section 3 we will discuss the issue of taxes and subsidies. In Section 4 we will discuss

[7] This point relates to the observation by Wittman (1984) that carrots induce participation while sticks discourage it. He did not consider, however, the interplay between compliance, violations and participation, which is the focus of our analysis.

[8] In the economic literature on tort liability, the participation decision is considered under the heading of "activity level" but the emphasis is on a different set of problems. The main concern of that literature is that it is impossible to incentivize optimal care (that is, compliance) and activity-level (that is, participation) decisions by two parties through ordinary liability rules (Shavell 1980). Here we are concerned with the behavior of a single agent in a population of agents and the main problem is that the policy is not tailored to the agent-specific characteristics. The literature on torts has only paid limited attention to this problem and only with respect to the cost of care (Ganuza and Gomez 2006).

[9] Also Levmore (1986) considered the joint use of carrots and sticks; he observed that a regulator trying to incentivize behavior with carrots may need to introduce sticks to prevent false claiming of carrots. That is precisely the problem we study here. Levmore's (1986) observation was brief, however, and did not analyze any heterogeneity among the agents.

the implications of the model for real-life regulatory settings and, in Section 5, we will conclude.

2 The model

2.1 Model setup

Consider a principal regulating an activity in which any of the agents in a population can elect to participate. Each agent chooses independently whether to participate in the activity and, if so, whether to comply with the regulation. (Those who do not participate are not subject to the regulation.)[10] Agents derive a fixed benefit from participation $b > 0$ and face a fixed cost of compliance $e > 0$.[11] Costs and benefits vary across agents, so that the population of agents can be fully described by the probability distribution $F(b,e)$ of the agents' two-dimensional types (b,e). Agents are risk neutral and maximize their expected payoff.

To induce agents to participate and comply, the principal monitors them with probability $0 < p \leq 1$,[12] and rewards compliance with a carrot $c > 0$ while punishing violations with a stick $s > 0$.[13] The principal's qualification of an agent's conduct as compliance or violation does not perfectly overlap with the agent's actual behavior. The principal correctly identifies compliance with probability $\frac{1}{2} < q_k < 1$ and, with the complementary probability $1 - q_k$, may erroneously classify an agent's conduct as violation while in fact the agent complied (a false positive, akin to convicting the innocent). Vice versa, the principal correctly identifies violation with probability $\frac{1}{2} < q_v < 1$ and, with probability $1 - q_v$, may erroneously detect compliance while the agent in fact did not comply (a false negative, akin to acquitting the guilty).[14]

2.2 Simplifying assumptions

To make our point in the starkest possible way, we simplify the general setup above along two dimensions. First, we make a number of simplifying assumptions about the probability distribution $F(b,e)$ of agents' types and characterize the population of agents as follows. We consider a population divided into two groups: a portion h of the agents are high-benefit agents and derive a benefit b_H from participation, the remaining portion $1 - h$ of the agents are low-benefit agents and derive a benefit b_L from participation, with

$$b_H > p(qs - (1-q)c) \geq b_L \tag{1}$$

[10] In Section 3, we show that our results also apply when the principal directly regulates participation through taxes and subsidies that are not conditional on agents' compliance.

[11] Non-participants earn no benefit and violators incur no cost. Some activities are costly to the agent; that is, we could allow b to be negative. Setting a zero lower bound for b, however, is without loss of generality. Allowing for negative benefits for participants is equivalent to setting a positive benefit for non-participants, which would not affect our results.

[12] As it will be clear from the setup, we focus here on an auditing model in which only monitored agents are eligible for the carrot and subject to the stick. The model however is robust to relaxations of this assumption, including offering the carrot to all participants unless they violate or similar variations. In the general model presented in the Appendix we will dispense with any restriction as to the carrot or stick applied to non-monitored agents.

[13] We define carrots and sticks in an unambiguous way: a carrot is a monetary payment from the principal to the agent conditional on (detected) compliance, while a stick is a monetary payment from the agent to the principal conditional on (detected) violation.

[14] By imposing that both probabilities are greater than $\frac{1}{2}$ we simply restrict attention to the plausible scenarios where the principal's enforcement technology fares better than a coin toss.

in the baseline setting. This assumption, as will be clarified below, guarantees that, in the baseline, all high-benefit agents choose to participate in the regulated activity, while low-benefit agents participate only if their cost of effort is low enough. This allows us to describe the behavior of agents in response to a policy change compared to the baseline in a simple way. Furthermore, we assume that, in both groups, the cost of compliance e is uniformly distributed between 0 and 1.[15]

Although the assumption in (1) will be relaxed in the Appendix, note that it is quite realistic. If the first inequality does not hold, the benefit of high-benefit agents is so low that they earn a negative payoff from participating and violating the rule and hence all participating agents would comply. Alternatively, if the second inequality does not hold, the benefit of low-benefit agents is so high that even non-complying agents participate. In this case, every single agent participates. Thus, only if (1) holds does the model reflect a realistic scenario, where some agents participate and comply, some agents participate and violate, and some agents stay out of an optional, imperfectly enforced regime.

The second set of assumptions we make relates to the accuracy of the principal: we assume that

$$q_k = q_v = q$$

that is, that the principal is equally accurate in detecting compliance and violations.[16] In the Appendix, we will relax all of these simplifying assumptions, show that the results remain unchanged and demonstrate some additional results, which are also discussed in what follows.

2.3 Analysis

Given the general setup in Section 2.1 and the simplifying assumptions in Section 2.2, an agent of type (b, e) anticipates the payoffs reported in Table 3.

Action	Payoff
The agent participates and complies	$\Pi_k (b, e) = b + pqc - p (1 - q) s - e$
The agent participates but violates	$\Pi_v (b) = b + p (1 - q) c - pqs$
The agent does not participate	0

Table 3: Agent's payoffs

Starting with the bottom row, if the agent does not participate, his (normalized) payoff is zero. If he participates but violates the principal's rule of conduct, the agent will earn the benefit b, will face a stick s with probability q (that is, if the principal correctly qualifies the agent's conduct as a violation of the rule), and will face a carrot with probability $1 - q$ (that is, if the principal erroneously qualifies the agent's conduct as compliance). The agent will incur no cost of effort. We denote the payoff of an agent who participates but violates the rule as $\Pi_v (b)$. Finally, if the agent participates and complies, he earns the benefit b, faces a carrot with probability q (that is, if the principal correctly qualifies his conduct as compliance), faces a stick with probability $1 - q$ (that is, if the principal erroneously qualifies his conduct as a violation), and incurs costly effort equal to e. We denote the payoff of an agent who participates and complies as $\Pi_k (b, e)$.

We can now characterize the agents' behavior, starting from high-benefit agents. Assuming that an agent participates, he will comply if the payoff from compliance is greater than the payoff

[15] Note that we impose independence between b and e. This will not be the case in the Appendix.

[16] It follows from the general setup that $\frac{1}{2} < q < 1$.

from violation, that is, if $\Pi_k(b_H, e) > \Pi_v(b_H)$, which is the case if the agent's cost of effort is lower than the compliance threshold

$$e_k \equiv p(c+s)(2q-1) \tag{2}$$

Note that, given the agent's choice between compliance and violation, the payoff of complying agents must be at least as high as the payoff of violating agents. Therefore, the lowest payoff a participating agent may earn is $\Pi_v(b_H)$. Given the assumption in (1), the payoff of high-benefit agents who violate is positive, $\Pi_v(b_H) > 0$, and hence all high-benefit agents participate in the activity irrespective of whether they comply or violate. Summing up, a portion e_k of the high-benefit agents participate and comply, while the remaining portion $1-e_k$ participate and violate.[17]

Turning now to the low-benefit agents, note that, by assumption (1), we have $\Pi_v(b_L) \leq 0$ and hence not all low-benefit agents participate in the activity. In particular, none of the would-be violators and some of the would-be compliers decide to opt out. Agents who participate must therefore be compliers and must earn a positive payoff, or $\Pi_k(b_L, e) > 0$, which is the case if the agent's cost of effort is lower than the participation threshold

$$e_p \equiv b_L + p(qc - (1-q)s) \tag{3}$$

To sum up, a portion e_p of the low-benefit agents participate and comply, while the remaining portion $1-e_p$ refrain from participation. Figure 2 below offers a visualization of the agents' behavior in a typical baseline setting, that is, when the assumption (1) holds.

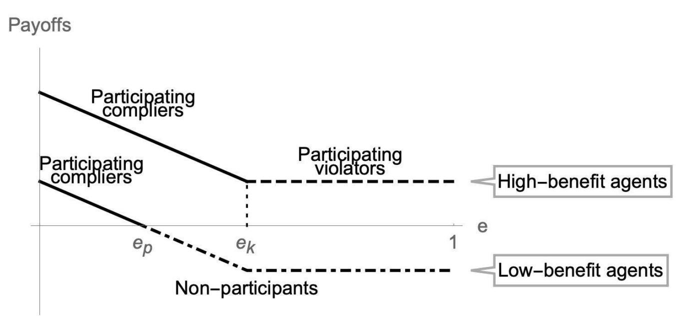

Figure 2: Agents' behavior (Simulation values: $c = s = 1$, $p = \frac{1}{2}$, $q = \frac{3}{4}$, $b_H = \frac{1}{2}$, $b_L = 0$)

2.4 Results

The traditional approach in the literature on the economic analysis of law enforcement rests on the implicit assumption that all agents (have to) participate in the regulated activity. In our simplified model, this approach is equivalent to focusing on the behavior of high-benefit agents. Those agents' choices are controlled by the compliance threshold e_k in (2), which increases in the carrot c, the stick s, and the monitoring probability p. It follows that, a marginal increase in any of these variables results in higher levels of compliance and, correspondingly, lower levels of violation. A number of further observations follow from the standard approach:

[17] This is due to the assumption that e is uniformly distributed between 0 and 1 so that the share of agents with $e < e_k$ is exactly e_k and, conversely, the share of agents with $e \geq e_k$ is $1 - e_k$.

1. Compliance and violations are complementary outcomes: an increase in compliance is mirrored by a decrease in violations and vice versa;

2. Carrots and sticks are substitutes: as long as the sum $c + s$ stays constant, any combination of c and s results in the same level of e_k and hence in the same level of compliance and violations;

3. The probability and the magnitude of sanctions are substitutes: as long as the product $p\,(c + s)$ (the expected sanction) stays constant, any combination of the probability p and the magnitude $(c + s)$ of the sanctions results in the same level of e_k and hence in the same level of compliance and violation.[18]

However, whenever agents have a meaningful choice whether to participate or to stay out of the regulated activity, the analysis needs to take into account also the behavior of the low-benefit agents, whose choices are controlled by the participation threshold e_p in (3). This threshold reacts differently to changes in the policy variables. In particular, the participation threshold e_p increases in the carrot c but decreases in the stick s. Moreover, e_p increases in the probability of monitoring p only if the carrot is large enough compared to the stick, and decreases in p otherwise. More precisely, e_p increases in p if $\frac{c}{s} > \frac{1-q}{q}$ and decreases in p if $\frac{c}{s} > \frac{1-q}{q}$. (Note that the latter condition is satisfied when $c = s$ because we have $\frac{1-q}{q} < 1$; in turn, this is because our assumption $q > \frac{1}{2}$ implies $q > 1 - q$).

For the low-benefit agents, the alternative to compliance is non-participation rather than violation and hence an increase in compliance does not reduce violations, rather, it induces broader participation. An increase in the carrots results in an increase in compliance and in participation, while an increase in the stick results in a decrease in compliance and participation. Finally, an increase in the probability of monitoring may results in an increase or in a decrease in compliance and participation.

These observations overturn the traditional understanding of agents' behavior in a regulatory framework in several fundamental ways. We start here by defining the baseline setting and then examine the effects of a change in c, s or p. To start, let us focus on a case in which the regulator employs a stick $s = 1$ and a carrot $c = 1$, monitors agents with probability $p = \frac{1}{2}$, and has accuracy $q = \frac{3}{4}$. The two groups of agents derive benefits $b_H = \frac{3}{2}$ and $b_L = \frac{1}{4}$, respectively. In this case, the compliance threshold is $e_k = \frac{1}{2}$ so that half of the high-benefit agents comply and the remaining half violate the rule (all of them participate); similarly, the participation threshold is $e_p = \frac{1}{2}$, so that half of the low-benefit agents participate and comply, while the remaining half opt out. Figure 3 illustrates this simple baseline setting. Below, we illustrate our results by departing from the baseline one variable at a time.

2.4.1 Compliance and violations are not complementary outcomes

A possibly relevant portion of the population may elect not to participate in the regulated activity and, as policy variables change, changes in compliance will not typically correspond to opposite changes in violations. In our example in Figure 3, the low-benefit agents at the margin choose between non-participation and compliance and a policy-induced reduction in compliance would result in less participation, not in more violations.

[18] Setting $c = 0$ results in the familiar formula for expected penalty, ps.

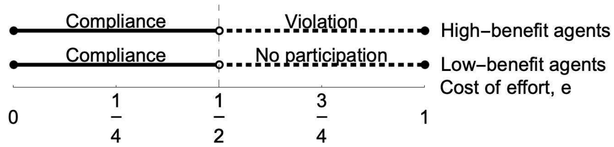

Figure 3: Baseline (Simulation values: $s = 1$, $c = 1$, $p = \frac{1}{2}$, $q = \frac{3}{4}$, $b_H = \frac{3}{2}$, $b_L = \frac{1}{4}$).

2.4.2 Increasing the carrot may result in an increase in violations

While the high-benefit agents react to an increase in carrots with a reduction in violations, the same increase in the carrots may push non-participating low-benefits agents to participate and hence will increase violations in that group from zero to some positive number. If there are enough low-benefit agents in the population, this effect will dominate the reduction in violations in the high-benefit group. Consider the example in Figure 3 and, all other things equal, increase the carrot from $c = 1$ to $c = 2$. This policy will increase the compliance threshold from $e_k = \frac{1}{2}$ to $e_k = \frac{3}{4}$, so that the high-benefit agents comply more and violate less, as shown in Figure [19]. However, the higher carrot also increases the payoff all low-benefit agents above 0, so that now they all participate.[20] Some of the former non-participants now participate and comply, while the remaining participate and violate the law, as shown in Figure [21] as compared to Figure 3. The total effect depends on the relative size of the two groups, that is, on h. If there are relatively many low-benefit agents (low h), the increase in violations in this group will more than offset the decrease in violations among high-benefit agents and an increase in the carrot will result in an increase in violations, contrary to the predictions of the standard approach.

Figure 4: Increased carrot (with $s = 2$, $c = 1$, $p = \frac{1}{2}$, $q = \frac{3}{4}$, $b_H = \frac{3}{2}$, $b_L = \frac{3}{4}$).

2.4.3 Increasing the stick may result in a decrease in compliance

An increase in the stick induces some complying low-benefit agents to opt out of the regulated activity. If there are enough low-benefit agents in the population, the decreases in participation of complying agents will dominate the increase in compliance among the high-benefit agents. Considering again the baseline described above, we now increase the stick from 1 to 2, which in turn results in an increase in e_k from $\frac{1}{2}$ to $\frac{3}{4}$ (as in the previous section) and a decrease in e_p from $\frac{1}{2}$ to $\frac{3}{8}$, as illustrated in Figure 5. While high-benefit agents comply more, low-benefit agents

[19] 4

[20] This is easy to see by noting that a violating low-benefit agent now earns $\Pi_v(b_L) = \frac{7}{8} > 0$ and hence participates. All complying low-benefit agents must earn at least that amount.

[21] 4

comply less and—using again the now familiar argument—if h is low enough, the total effect of an increase in the stick will be a reduction in compliance, again contrary to the predictions of the standard approach.

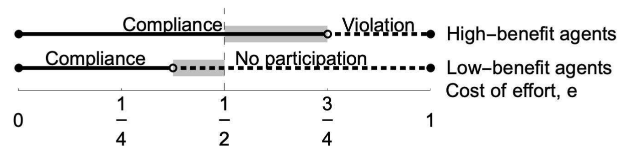

Figure 5: Increased stick (with $s = 1$, $c = 2$, $p = \frac{1}{2}$, $q = \frac{3}{4}$, $b_H = \frac{3}{2}$, $b_L = \frac{3}{4}$).

2.4.4 Carrots and sticks are not (perfect) substitutes

Combining the previous two results, it is now easy to see that carrots and sticks are substitutes only if one focuses exclusively on the high-benefit agents, who always participate. For the low-benefit agents, carrots are sticks are (imperfect) complements because an increase in carrots results in outcomes that could be alternatively achieved by a reduction in the stick. Hence, in the aggregate, an increase in carrots and an increase in sticks will typically results in different outcomes. More starkly, they may result in opposite outcomes.

In the examples above, we have increased either the carrot or the stick by the same amount; that is, we have added $\frac{1}{2}$ to the relevant status-quo level in both cases. The behavior of the high-benefit agents changes in exactly the same way whether we increase the carrot or the stick; which replicates the standard approach. However, when the behavior of low-benefit agents is also accounted for, results change. While increasing the carrot unambiguously increases compliance, we have shown that increasing the stick may result in a reduction in compliance. Similarly, while an increase in the stick reduces violations, an increase in the carrot may result in an increase in violations.

2.4.5 Increasing the probability of monitoring may increase or decrease participation, compliance and violations

Only sticks or only carrots If only carrots or only sticks are used, then what counts is the expected carrot or stick and hence the effect of increasing the probability of monitoring is the same as the effect of increasing the magnitude of the corresponding sanction—that is, it is it is the same as increasing the carrot if only carrots are used and it is the same as increasing the stick if only sticks are used. This includes our result about the ambiguous effect of carrots and sticks: increasing the probability of monitoring has an ambiguous effect on compliance if only sticks are used and an ambiguous effect of violations if only carrots are used. In this simple setup, the magnitude and the probability of sanctions are perfect substitutes. Things are more interesting when carrots and sticks are used in tandem, as explained in the following.

Both carrots and sticks In the standard approach, incentives are produced by the expected sanction—that is, $p(s+c)$ or, if only sticks were used as in our baseline example, ps—and hence a 50% increase in the probability of monitoring and a 50% increase in the magnitude of sanctions produce the same levels of compliance and violations. This is indeed the case for high-benefit agents, but may or may not be the case for low-benefit agents. The effect of a change in the probability of monitoring for the low-benefit agents depends on the relative size of carrots and sticks—as can be seen from (3)—and hence any outcome is possible if there are enough low-benefit agents in the population.

Let us consider again our baseline example and note that in the analysis above we increased the magnitude of the sanctions, $c + s$, by 50% by increasing either the carrot or the stick from 1 to 2. In both cases the total magnitude of the sanctions $c + s$ increased from 2 to 3, a 50% increase. Here we consider an analogous 50% increase in the probability of monitoring from $p = \frac{1}{2}$ to $p = \frac{3}{4}$ and examine its effect on the behavior of the low-benefit agents. (High-benefit agents react by complying more and violating less, as we know.) Again, we have an increase in e_k from $\frac{1}{2}$ to $\frac{3}{4}$ and as increase in e_p from $\frac{1}{2}$ to $\frac{5}{8}$, as illustrated in Figure 6. Note that the effect is different from both the effect of increasing the carrot and that of increasing the stick. The sign of the change is similar to that of a carrot (there is more compliance) because carrots are relatively large. The opposite would be true if sticks were higher relative to carrots.

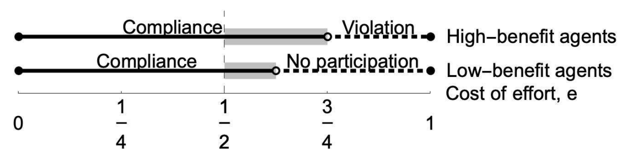

Figure 6: Agents' behavior (with $s = 2$, $c = 0$, $p = \frac{3}{4}$, $q = \frac{3}{4}$, $b_H = \frac{3}{2}$, $b_L = \frac{3}{4}$).

2.4.6 The probability and magnitude of sanctions are not substitutes

For the low-benefit agents, increasing the probability of monitoring has the same effect as increasing the carrot if carrots are large enough and has the same effect as increasing the stick in the opposite case. Thus, increasing the probability of monitoring is either a substitute for carrots or a substitute for sticks, but never of both. In the aggregate, an increase in the probability of sanctions will typically result in different outcomes when compared with an increase in the magnitude of sanctions.

Depending on the relative ratios of high- and low-benefit agents in the population—that is, depending on h—the outcome of basic policy interventions on the magnitude and probability of sanctions may be dramatically different from what would be predicted by the standard approach.

3 Taxes and subsidies

Let us now examine whether adding a tax to the model expands the set of outcomes that the principal can reach. Assume that the principal taxes non-participation in the activity by a fixed amount τ. Obviously, and in contrast with the carrot and the stick in our model, this tax does not depend on the compliance decision of participating agents. Alternatively, the principal subsidizes the regulated activity (again, regardless of the principal's findings of compliance or violation) by the same amount τ. (If τ is negative we have a tax for participants or, equivalently, a subsidy for non-participants.) The new payoff matrix is as follows:

Action	Payoff
The agent participates and complies	$\Pi_k(b,e) = b + pqc - p(1-q)s - e + \tau$
The agent participates but violates	$\Pi_v(b) = b + p(1-q)c - pqs + \tau$
The agent does not participate	0

Table 4: Agent's payoffs

We have now the following thresholds:

$$\begin{aligned} e_k^{\tau} &= (c+s)\,p\,(2q-1) \\ e_p^{\tau} &= b_L + p\,(qc - (1-q)\,s) + \tau \end{aligned}$$

Clearly the compliance threshold is unaffected by τ, because every participant receives τ whether he complies or not. In contrast, participation will be incentivized by τ, both for compliers and for violators, since increases $e_p^{\tau}(b)$. The addition of τ, however, does not expand the set of outcomes that the principal can reach. To see why, consider that the principal could replicate the same compliance and participation thresholds as above by simply adjusting the carrot and the stick instead of implementing a tax or a subsidy. More precisely, the principal could increase the carrot by an amount equal to $\frac{\tau}{p}$ and reduce the stick by the same amount; that is, the principal could set $c' = c + \frac{\tau}{p}$ and $s' = s - \frac{\tau}{p}$ to mimic the thresholds reported above. The compliance threshold is clearly met, because $c' + s' = c + s$. The participation threshold is also met because we have

$$\begin{aligned} e_p(b) &= b + p\left(qc' - (1-q)\,s'\right) \\ &= b + p\left(q\left(c + \frac{\tau}{p}\right) - (1-q)\left(s - \frac{\tau}{p}\right)\right) \\ &= b + p\,(qc - (1-q)\,s) + \tau \\ &= e_p^{\tau}(b) \end{aligned}$$

In essence, the principal can replicate the effects of taxes and subsidies by keeping the sum of the carrot and the stick constant (so as to keep the compliance threshold constant) while altering their ratio $\frac{c}{s}$. That ratio should increase if the activity is subsidized (that is, if τ is positive) and decrease if the activity is taxed (that is, if τ is negative).

4 Discussion

Our results upend the long-held views about the basic workings of deterrence mechanisms and the interplay between positive and negative incentives. Yet our analysis hews closely to Becker's (1968) canonical model. We modify that model in only three ways. First, we add the participation margin to the analysis of compliance. Second, we add carrots to the law enforcement setting. And third, we introduce agents who differ in at least one dimension in addition to their varying costs of complying with the law. In this section, we explain that our modifications are neither minor embellishments nor artificial contrivances. Rather, these modifications make the law enforcement model more realistic and more useful to real-world regulators.

4.1 Optional participation

Consider some of the canonical examples of socially costly acts in the law and economics literature: pollution (Coase 1960, Calabresi and Melamed 1972), speeding (Becker 1968, Stigler 1970), collusion and monopolies (Becker 1968, Stigler 1970, Tullock 1967), theft (Becker 1968, Calabresi and Melamed 1972, Tullock 1967), tax evasion (Allingham and Sandmo 1972). All these acts–as well as many others–raise the question of deterring harmful conduct. Most of the economic analysis of deterrence–also known as the theory of public enforcement of law–is preoccupied with answering this question.

In contrast, scholars of deterrence have paid little attention to the effect of enforcement instruments on the agent's decision to participate in a regulated activity. There is some work studying the extent of participation (the level of activity), mostly conditional on the optimally set deterrence incentives (Polinsky and Shavell 2007: 425). In addition, Png (1986) and Kaplow (2011a, 2011b) investigated the interaction between deterrence (compliance) and chilling (participation). Png (1986) offered initial analysis, and Kaplow (2011a, 2011b) focused on the implications for the burden of proof.

Yet as we emphasized in the Introduction, the decision to participate is of first-order importance in the regulation of pollution, speeding, collusion, theft, tax evasion, and so on. A firm potentially liable for pollution may exit the "dirty" industry; a firm potentially liable for collusion may exit the concentrated sector of the economy; a driver potentially liable for speeding may stay home or get on a bicycle, a taxpayer concerned with sanctions for non-compliance may avoid transactions with uncertain tax consequences; and so on. In fact, regulatory settings when the non-participation option is unavailable are the exception, not the rule.

Moreover, the concerns with potential liability just discussed—including potential mistaken liability—are real. Government regulation in all of the areas just mentioned (and in many others) is far from error-free. Vague standards are commonplace, and a mistaken imposition of liability is far from a remote possibility. Given this concern, opting out of the activity altogether is not a theoretical consideration but a very real choice. A model that ignores this choice necessarily has only a limited explanatory power when applied to the real world.

In fact, this limitation is even more significant than it has appeared so far. All examples considered until now involve potentially harmful acts that yield private benefits that may or may not exceed the external harms. There is, however, another category of acts-–those producing private benefits that are too small (compared to private costs) to induce individuals to undertake them. If the government wants these acts to occur, it needs so supply additional incentives-–carrots, in the law and economics parlance. Obviously, the question of what carrots are sufficient to induce participation in these explicitly optional regimes is central the their analysis. We discuss these regimes next.

4.2 Carrots, not just sticks

The traditional analysis of law enforcement posits the government that deters privately beneficial but socially undesirable acts by imposing sanctions equal to the acts' external harms in expectation. Thus the key variables are the private benefit, the external harm, and the expected sanction. Positive incentives appear to have no role in this setup.

Png (1986) suggested that carrots should be used to counter the erroneous imposition of sticks on complying agents. But Png's (1986) brief discussion did not address the practical significance of carrots in law enforcement. His example was wildly unrealistic,[22] and his insight has remained undeveloped. At the same time, a voluminous literature on carrots and sticks focuses on their effectiveness in inducing participation in a particular activity. That literature generally does not address the problem of improper (illegal) claiming of carrots. In sum, the deterrence scholarship focuses on sticks and enforcement errors while mostly ignoring carrots, and the literature on carrots and sticks mostly ignores the imperfect enforcement.

Yet a moment's reflection leads to two conclusion. First, optional incentive-based regimes do raise noncompliance concerns. Second, traditional enforcement settings studied in deterrence literature do feature carrots, though usually of the unintended, quasi-carrot variety.

Starting with the optional, incentive-based regime, note that the US government spends multiple billions each year on cash grants (Goodwin and Smith 2018; Theodos, Stacy and Ho 2017), loan assistance (Goodwin and Smith 2018), price and revenue protection programs (Glauber 2018; Sumner 2018), insurance subsidies (Glauber 2018; Sumner 2018), and other carrots designed to induce participation in a wide variety of government schemes aimed at anything from peanut farming (Goodwin and Smith 2018) to community development (Theodos, Stacy and Ho 2017) to workplace diversity (Sturm 2006). Small businesses alone benefit from fifty two different federal assistance programs administered by the Departments of Agriculture, Commerce, Housing and Urban Development, and the Small Business Administration (U.S. Government Accountability Office 2018). U.S. states spend billions more on similar inducements (Patrick 2014).

Some of the carrot-granting programs challenge the recipients (and the regulators) with "a complicated labyrinth of criteria", making enforcement errors all but certain (Howells 19970, pp. 393-94, 412). Some carrot-claimers skirt the rules, failing to comply either with their letter or their spirit (Goodwin and Smith 2018; Howells 1997; Theodos, Stacy and Ho 2017). In response, regulators use penalties to deter improper carrot-claiming (Howells 1997). Thus simultaneous presence of carrots, sticks, and enforcement errors is a typical feature of many US incentive programs, not a rare aberration.

The U.S. government also spends billions more in positive incentives outside of the federal budget. The so-called tax expenditures—special tax breaks in the Internal Revenue Code designed to induce participation in a particular activity—offer potential participants a carrot just like outright cash grants do. And for the largest eight hundred or so companies under a continuous IRS audit (Elliott 2012), the only uncertainty surrounding the carrot arises from inaccurate enforcement by the government auditors that is key to our model.

In fact, all three of the largest corporate tax expenditures of the recent past—the deduction for accelerated depreciation, the domestic production activities deduction (repealed at the end of 2017), and the research and development (R&D) tax credit—illustrate every feature of our model. The R&D credit, to take one example, is a carrot. Its false claiming is punished with a tax penalty stick. Over the years, Congress has varied both on multiple occasions. The same is true of the accuracy of the government's eligibility determinations. Given that Congress tends to tinker with tax incentives particularly often, our analysis should be especially valuable in that context.

Two further insights highlight the wide application of our results. First, the carrot need not take the form of a payment or a reduction in a price or tax liability. Any government-created benefit is a carrot that may be a part of a selective incentive regime. The U.S. patent system is a case in point. The government incentivizes innovation with a carrot of exclusive use of the patented invention during the patent term. The invention's novelty is an imperfectly observable

[22] Png suggested a carrot in the form of payments to all drivers involved in accidents, whether they exercise due care or not, Png (1986: 103).

characteristic that occasionally leads to inaccurate patent-granting decisions by the U.S. Patent and Trademark Office. And Congress has tried to deter bad-faith claims on novelty by penalizing them with criminal fines—a stick if there ever was one (Yelderman 2017).

The second insight is that the carrot need not be intended as such by the regulator. Consider the Federal Reserve's so-called discount window that offers eligible banks access to short-term liquidity under certain conditions. The discount window was put in place in 1913 to establish the Fed as the lender of the last resort, and possibly to incentivize small-business lending by local banks (Hunt 2014). It was certainly not created to incentivize trading firms to become banks. Yet this was exactly how it worked when Morgan Stanley and Goldman Sachs chose to become banks in the midst of the financial crisis (Sorkin and Bajaj 2008). The discount window acted as an unintended quasi-carrot.

The Volcker Rule enacted in the aftermath of the crisis in order to stop proprietary trading by deposit-taking institutions completed the quasi-incentive regime that we model. The discount window is a quasi-carrot. The broad enforcement powers of the Fed are a stick (Hamilton 2017). And the purpose-based definition of proprietary trading, as well as the hedging and market-making exceptions, in the Volcker Rule inevitably give rise to false positives and false negatives.

Tax quasi-carrots are the focus of the vast amount of day-to-day corporate tax planning. The tax-free reorganization provisions of the Internal Revenue Code were not enacted to induce tax-free reorganizations. Rather, Congress put them in place under "the assumption [ ] that the new enterprise ... that may hold the corporate assets, and the new stock or securities received in exchange for old stock or securities, are substantially continuations of, and interests in, the old corporations, still alive but in different form" (Bittker and Eustice Par. 12.00[1] 2019). Yet, there is no doubt that massive savings from qualifying corporate acquisitions and separations for a tax-free treatment are of first-order significance in momentous business decisions that shape financial markets and the US economy (Elliott 2015). And the voluminous tax literature exploring whether the tax-free treatment is available because transactions take place pursuant to a "plan," whether "substantially all" of the assets are transferred, and whether the entire reorganization lacks "economic substance," among other conditions, makes it quite clear that their application involves plenty of uncertainty.

Likewise, Congress enacted the US international tax rules—including the foreign tax credits and the reduced rate on active foreign earnings—in order to prevent double-taxation of income earned by US companies. The goal was certainly not to induce US companies to move their operations or corporate domicile offshore. In fact, the major reform of US international tax regime in 2017 was motivated by the desire to prevent such shifts, including the so-called corporate inversions. Yet, it is beyond doubt that foreign tax credits and reduced corporate tax rates do serve as a quasi-carrot-–an unintended inducement to move operations, headquarters, and income offshore. At the same time, the complexity of the international tax rules makes it impossible to enforce them error-free.

Overall, given the multitude of explicitly optional incentive-based schemes, and the numerous actually optional settings featuring quasi-carrots, our model featuring both sticks and carrots in the presence of imperfect enforcement reflects a very large number of real-world regulatory settings.

4.3 Heterogeneous agents

Our last modification of the standard deterrence model is the introduction of agents who vary along a dimension other than their cost of compliance. Thus far, that dimension was the agents' private benefit from participation. One hardly needs to spend much time to defend this assumption. In fact, the opposite assumption would be wildly unrealistic. Multiple polluting plants do not all derive the same benefits from running their operations. Multiple drivers do not benefit equally from getting to their destinations. Multiple reorganizing firms-–even while undertaking the same type of a tax-free reorganization—do not derive the same benefit from their business combinations. Private benefits are, indeed, heterogeneous.

Our results require that agents vary in one dimension other than their costs of compliance, but that dimension is not limited to private benefits. Another way in which agents can (and, no doubt, do) vary is their perception of enforcement accuracy. This variation is so likely that assuming it away would considerably limit the realism of the model. Some agents perceive the principal to be highly accurate; others believe that the principal is rather error-prone. In the vast majority of regulatory setting there is simply no mechanism for agents to share their impressions of the government's enforcement accuracy and to reach a universally-agreed-upon view. We do not assert that agents have no information about the matter. Litigation, published official interpretations, and public statements of government officials at industry gatherings all provide agents with valuable information. But it is implausible to assume that given that information all agents arrive at the same conclusion regarding the government's enforcement error rates. As we show, heterogeneity of perceptions of the government's enforcement accuracy yields the same key results as the heterogeneity of benefits does.

In the model considered thus far, agents vary with respect to their costs of effort and the benefit that they derive from participation. However, as we show in detail in the Appendix, our results hold also in a setting in which agent derive the same benefit from participation but hold different beliefs as to the accuracy of the principal in detecting compliance and violation, that is, agents vary in their perceptions of q_k and q_v. This may be for two reasons. First, whatever principal believes about her own accuracy, it is difficult for the principal to convey this belief to the agents, and all but impossible to do so credibly. Second, it is implausible to assume that the principal is equally accurate in evaluating compliance of every agent. This suggests that our results are not limited to the particular setup we present in the main part of the paper, but hold more generally when the agent's decision to participate is affected by some agent-specific characteristic other then the agent's cost of compliance with the regulation.

4.4 When the conventional wisdom is wrong

...to be added...

5 Conclusion

Inevitably, our analysis has limitations. We discuss five of them here. The first limitation reflects a well-known and persistent challenge to the economic analysis of law. Agents in our model cannot affect the principal's accuracy by adjusting their behavior. Agents either comply or violate; they cannot barely cross the line or engage in egregious noncompliance. This assumption is more realistic in some settings than in others.

The assumption of agents' binary behavior is standard in law and economics. The reason for it is well-known: eliminating this assumption leads to indeterminate results (Craswell and Calfee 1986; Shavell 1987). Several recent contributions show that in certain settings this indeterminacy disappears (Baker and Raskolnikov 2017; Dari-Mattiacci 2005). Our analysis does not contribute to further resolving this long-standing challenge.

The second limitation is our principal's inability to tailor its instruments to individual agents. This inability is one of the key distinctions between a constrained model like ours and the mechanism design literature (Kaplow 2017). In reality, regulators usually deploy crude, uniform instruments rather than variable incentives (Kaplow 2017, Montero 2005, Stavins 2003). So our assumption of uniform carrots, sticks, and likelihood of detection is quite realistic in many settings.

Third, participants in our model have a choice of complying or violating the rule, but they cannot choose the extent of participation. Incorporating this variable into the model would necessarily complicate it. Our goal is to highlight that extending the standard model to reflect a small number of realistic features (participation decision, presence of carrots, heterogeneity of agents) is sufficient to modify the canonical conclusions of the deterrence theory.

Fourth, we do not model costs of various regulatory tools that we consider because we are not concerned with welfare-maximizing policy interventions. These costs are likely to vary. Carrots are usually more costly for the principal than sticks. Increasing the rate of monitoring is likely to consume more of the principal's resources than adjusting a monetary sanction. Our contribution is to clarify the effects of changing different enforcement instruments, leaving it to the principals (and to future work) to balance these effects against the costs.

Finally, and at the risk of stating the obvious, our model is a simplified version of reality. Not only agents' benefits and principal's accuracy vary simultaneously in real-world settings, other factors such as risk aversion, uncertainty aversion, and various heuristics and biases likely come into play.

Our decision to ignore all these complications is deliberate. We re-examine the foundational conclusions of a canonical model of law enforcement, and we show that these conclusions do not generally hold even without adding any risk aversion, behavioral factors, or other deviations from rationality. Extending the model to incorporate such deviations is likely to raise further questions about general applicability of the model's basic payoffs.

In the end, our unexpected findings as well as all of the limitations just discussed sound a note of caution for real-world regulators. Rather than relying on the seemingly straightforward prescriptions of the basic deterrence model, they would do well by closely monitoring the actual responses to any change in actual selective incentives (Duflo 2017).

References

[1] Andreoni, James, William Harbaugh, and Lise Vesterlund. 2003. The Carrot or the Stick: Rewards, penalties, and Cooperation, 93 *American Economic Review* 93:893-902.

[2] Baker, Scott, and Alex Raskolnikov. 2017. Harmful, Harmless, and Beneficial Uncertainty in Law. *Journal of Legal Studies* 46: 281-307.

[3] Becker, Gary S. 1968. Crime and Punishment: An Economic Approach. Journal of Political Economy 76: 169-217.

[4] Ben-Shahar, Omri, and Anu Bradford. 2012. Reversible Rewards. *American Law and Economics Review* 15:156-80.

[5] Bernstein, Lisa. 2015. Beyond Relational Contracts: Social Capital and Network Governance in Procurement Contracts. *Journal of Legal Analysis* 7:561-621.

[6] Bruch, Nicholas, Eric Dewey, and Steven Kovalan. 2015. Surmounting the Lateral Partner Hiring Challenge: Lessons Learned, Best Practices, and Tools for Success. *ALM Legal Intelligence*.

[7] Cecchi-Dimeglio, Paola, and Hugh A. Simons. 2018. Lateral Moves: An Empirical Investigation of Cyclicality, Directional Mobility, and 5-Year Retention Rate by Gender and Age Cohort. *Journal of Legal Profession* 42:171-195.

[8] Dari-Mattiacci, Giuseppe. 2005. Errors and the Functioning of Tort Liability. *Supreme Court Economic Review* 13:165–87.

[9] Dari-Mattiacci, Giuseppe. 2009. Negative Liability, *Journal of Legal Studies* 38:21-60.

[10] Dari-Mattiacci, Giuseppe, Gerrit De Geest and Jacques J. Siegers. 2009. Annullable Bonuses and Penalties. *International Review of Law and Economics* 29:349-359.

[11] De Geest, Gerrit and Giuseppe Dari-Mattiacci. 2013. The Rise of Carrots and the Decline of Sticks. *University of Chicago Law Review* 80:341-93.

[12] Duflo, Esther. 2017. The Economist as Plumber. *American Economic Review: Papers & Proceedings* 107: 1-26.

[13] Easterbrook, Frank H. 1985. Detrebling Antitrust Damages. *Journal of Law and Economics* 28: 445-468

[14] Elliott, Amy. 2012. Audit Proof? How Hedge Funds, PE Funds, and PTPs Escape the IRS. *Tax Notes* 136: 351-355.

[15] Emons, Winand. 1994. The Provision of Environmental Protection Measures Under Incomplete Information: An Introduction to the Theory of Mechanism Design. *International Review of Law and Economics* 14:479-91.

[16] Galanter Marc, and William Henderson. 2008. The Elastic Tournament: A Second Transformation of the Big Law Firm. *Stanford Law Review* 60:1867-1929.

[17] Galle, Brian. 2012. The Tragedy of the Carrots: Economics and Politics in the Choice of Price Instruments, *Stanford Law Review* 64:797—850.

[18] Organza, Juan-Jose, and Fernando Gomez. 2006. Caution, Children Crossing: Heterogeneity of Victim's Cost of Care and the Negligence Rule. *Review of Law and Economics* 1: 365–397.

[19] Gilpatric, Scott M. 2009. Risk Taking in Contests and the Role of Carrots and Sticks, *Economic Inquiry* 47:266-77.

[20] Gilson, Ronald J., Charles F. Sabel, and Robert E. Scott. 2018 The Adaptive Contract: Collaboration and Innovation in an Uncertain World. Unpublished manuscript. Columbia Law School.

[21] Glauber, Joseph A. 2018. Unraveling Reforms? Cotton in the 2018 Farm Bill. *American Enterprise Institute*.

[22] Gluckman, Nell. 2017. Changing Compensation Strategies Put Partners Under Pressure. *American Lawyer*, February 2017.

[23] Goodwin, Barry K., and Vincent H. Smith 2018. Reflections on the US Peanut Program. *American Enterprise Institute*.

[24] Graetz, Michael J., and Louis L. Wilde. 1985. The Economics of Tax Compliance: Fact and Fantasy. *National Tax Journal* 38:345-63.

[25] Hamilton, Jesse. 2017. Deutsche Bank Is First Bank Busted for Breaking Volcker Rule. *Bloomberg*, April 20.

[26] Henderson, William D. 2006. An Empirical Study of Single-Tier versus Two-Tier Partnerships in the AmLaw 200. *North Carolina Law Review* 84:1691-1750.

[27] Henderson, William D., and Leonard Bierman. 2009. An Empirical Analysis of Lateral Lawyer Trends from 2000 to 2007: The Emerging Equilibrium for Corporate Law Firms. *Georgetown Journal of Legal Ethics* 22:1395-1420.

[28] Henderson, William D., and Christopher Zorn. 2015. An Rx for Lateral Heartache. *American Lawyer*, February 9.

[29] Hildebrandt Consulting. 2018. Citi 2018 Client Advisory.

[30] Howells, Louise A. 1997. Looking for The Butterfly Effect: An Analysis of Urban Economic Development Under the Community Development Block Grant Program. *Saint Louis University Public Law Review* 16:383-417.

[31] Hunt, Daniel J. 2014. Just Grin and Bear It: Why Consistent Use of Individual Bailouts Under Section 13(3) of the Federal Reserve Act Is a Necessary Evil to Combat Economic "Mass Destruction." *George Mason Journal of International Law* 6:59-105.

[32] Kaplow, Louis. 1994. Accuracy in the Determination of Liability. *Journal of Law and Economics* 37:1-15.

[33] ____________. 2011. Optimal Proof Burdens, Deterrence, and the Chilling of Desirable Behavior. *American Economic Review: Papers & Proceedings* 101:277-80.

[34] ____________. 2012. Burden of Proof. Yale Law Journal 121:738-859

[35] ____________. 2017. Optimal Regulation with Exemptions. Working Paper No. 23887. National Bureau of Economic Research, Cambridge, Mass.

[36] Kenworthy, Lane, Stewart Macaulay, and Joel Rogers. 1996. "The More Things Change . . . ": Business Litigation and Governance in the American Automobile Industry. *Law and Social Inquiry* 21:631-78.

[37] Kleven, Henrik Jacobsen, and Wojciech Kopczuk. 2011. Transfer Program Complexity and the Take-Up of Social Benefits, *American Economic Journal: Economic Policy* 3:54-90.

[38] Lazear, Edward P. 1991. Labor Economics and the Psychology of Organization, *Journal of Economic Perspectives* 5:89-110.

[39] Levmore, Saul. 1986. Waiting for Rescue, An Essay on the Evolution and Incentive Structure of the Law of Affirmative Obligations. *Virginia Law Review* 72:879-941.

[40] Moldovanu, Benny, Aner Sela, and Xianwen Shi. 2012. Carrots and Sticks: Prizes and penalties in Contests, *Economic Inquiry* 50:453-62.

[41] Montero, Juan-Pablo. 2005. Pollution Markets with Imperfectly Observed Emissions. *RAND Journal of Economics* 36: 645-660.

[42] Olson, Mancur. 1992. Do Lawyers Impair Economic Growth. *Law and Social Inquiry* 17:625-633.

[43] Patrick, Carlieanne. 2014. Does Increasing Available Non-tax Economic Development Incentives Result in More Jobs? *National Tax Journal* 67:351-86.

[44] Png, I.P.L. 1986. Optimal Subsidies and Damages in the Presence of Judicial Error, *International Review of Law and Economics* 6:101-105.

[45] Polinsky, Mitchell A. 1979. Notes on the Symmetry of Taxes and Subsidies in Pollution Control. *Canadian Journal of Economics* 12:75-83.

[46] Polinsky, Mitchell A., and Steven Shavell. 2007. The Theory of Public Enforcement of Law. Pp 403-454 in vol. 1 of *The Handbook of Law and Economics*, edited by Mitchell Polinsky and Steven Shavell. Amsterdam: North Holland.

[47] Press, Aric, and Mary Cho. 2014. The State of the Big Law Market. *ALM Legal Intelligence Report*.

[48] Raskolnikov, Alex. 2014. Irredeemably Inefficient Acts: A Threat to Markets, Firms, and the Fisc. *Georgetown Law Journal* 102: 1133-1189.

[49] Shavell, Steven. 1980. Strict Liability versus Negligence. *Journal of Legal Studies* 9: 1-25.

[50] ____________. 2017. *******. Working Paper. Harvard University Law School. Cambridge, MA.

[51] Spector, Mike and Megumi Fujikawa. 2016. Millions More Cars With Takata Air Bags Recalled; Honda, Fiat Chrysler, Toyota, Nissan, Subaru and More Kick Off Latest U.S. Recalls. *Wall Street Journal*, May 27.

[52] Simons, Hugh A., and Michael A. Ellenhorn. 2018. Recent Survey Shows Firms Aren't Doing Right Due Diligence on Laterals. *American Lawyer*, January 28.

[53] Sorkin, Andrew Ross, and Vikas Bajaj. 2008. Shift for Goldman and Morgan Marks the End of an Era. *New York Times*, September 21.

[54] Stavins, Robert N. 2003. Experience with Market-Based Environmental Policy Instruments. Pp 355-435 in Vol. 3 of *Handbook of Environmental Economics*, edited by Karl-Göran Mäler and Jeffrey Vincent. Amsterdam: Elsevier Science.

[55] Sumner, Daniel A. 2018. Dairy Policy Progress. *American Enterprise Institute*.

[56] Theodos, Brett, Christina Plerhoples Stacy, and Helen Ho. 2017. Taking Stock of the Community Development Block Grant. Urban Institute. Washington D.C.

[57] U.S. Government Accountability Office. 2018. *Economic Development and Small Business*.

[58] Wald, Matthew. 2000. Tread Failures Lead to Recall of 6.5 Million Firestone Tires. *New York Times*, August 10.

[59] Weisbach, David. 2002. Ten Truths About Tax Shelters. *Tax Law Review* 55: 215-253.

[60] Wittman, Donald. 1984. Liability for Harm or Restitution for Benefit?, *Journal of Legal Studies* 13:57-80.

[61] Yelderman, Stephen. 2017. The Value of Accuracy in the Patent System. *The University of Chicago Law Review* 84:1217-1284.

A Appendix

A.1 General model

In this Appendix we prove that the results presented in the main text of the paper hold generally in the model setup presented in Section (2.1), that is, also if we relax all of the simplifying assumptions made in Section (2.2). In addition, we generalize these results and provide additional technical insights.

A.1.1 The model with perfect monitoring ($p = 1$).

We start by proving our results for the case in which the probability of monitoring is $p = 1$ and later extend the analysis to $p < 1$. Given the setup in Section (2.1), an agent of type (b, e) anticipates the payoffs reported below in Table 5, which follows the same logic as Table 3 presented in the main text.

Action	Payoff
The agent participates and complies	$\Pi_k (b, e) = b + q_k c - (1 - q_k) s - e$
The agent participates but violates	$\Pi_v (b) = b + (1 - q_v) c - q_v s$
The agent does not participate	0

Table 5: Agent's payoffs

It is useful to begin the analysis by defining the following threshold level for b

$$b_v \equiv q_v s - (1 - q_v) c$$

which is the level of b such that $\Pi_v (b) = 0$. Specifying this benefit threshold b_v allows us to partition the population of agents in two groups:

- High-benefit agents (with $b > b_v$) participate in the activity irrespective of whether they comply or violate. To see why, note that if $b > b_v$, then $\Pi_v (b) > 0$ and hence the payoff of violators is positive. If the agent complies, he must also earn a positive payoff because he will comply only if $\Pi_k (b, e) > \Pi_v (b)$, and we know that the latter is positive. Overall, high-benefit agents form the full participation group—they all take part in the activity. Whether the agent complies or violates depends on his cost of effort e. More precisely, the agent will comply if $e < e_k$ where the compliance threshold e_k is

$$e_k \equiv (c + s) (q_k + q_v - 1)$$

which is the level of e such that $\Pi_k (b, e) = \Pi_v (b)$. The agent will violate if $e \geq e_k$. When the compliance threshold controls, carrots and sticks are substitutes.

- Low-benefit agents (with $b \leq b_v$) participate in the activity only if they comply. To see why, note that if $b \leq b_v$, then $\Pi_v (b) \leq 0$ and hence the payoff of violators is (weakly) negative. The payoff for compliers may be positive or negative. Overall, low-benefit agents form the partial participation group where all participants comply, all violators abstain from participation, and some would-be compliers choose not to participate as well. A potential complier will participate only if $\Pi_k(b,e) > 0$ at his level of e. More precisely, the agent will comply if $e < e_p$ where the participation threshold e_p is

$$e_p(b) \equiv b + q_k c - (1 - q_k) s$$

which is the level of e such that $\Pi_k(b,e) = 0$. The agent will participate and comply if $e < e_p(b)$ and will not participate if $e \geq e_p(b)$. Note that $e_p(b)$ controls both compliance and participation because an agent who does not have incentives to comply would earn a negative payoff by participating and violating. When the participation threshold controls, carrots and sticks are complements.

In sum, e_k is the compliance threshold for participants, which determines the choice between complying and violating, on the assumption that the agent participates. This threshold controls the agent's compliance decision if $b > b_v$, that is, in cases where both complying and violating agents participate. In contrast, $e_p(b)$ is the participation threshold for complying agents. It controls behavior when $b \leq b_v$, that is, when violating agents do not find it advantageous to participate. In this case, the alternative to compliance is abstention rather than violation.

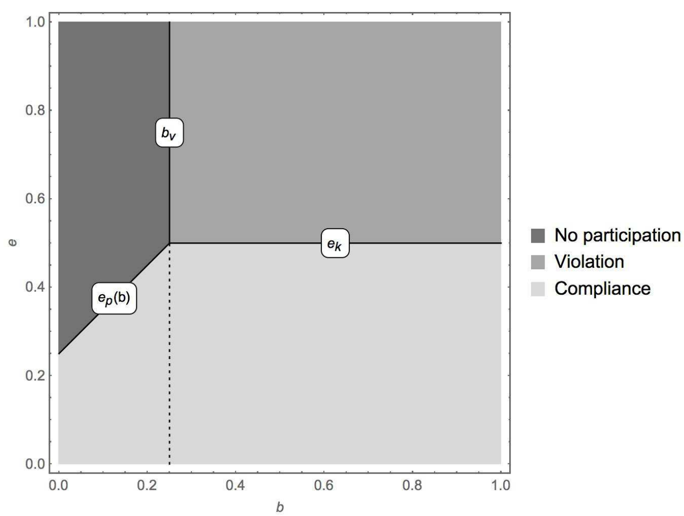

Figure 7: Agent's behavior (simulation parameters: $q_k = q_v = \frac{3}{4}$, $c = s = \frac{1}{2}$).

We impose no constraints on the possible benefits of heterogeneous agents in our model. Therefore, some agents in our model belong to the full participation group while others belong to the partial participation group. Figure 7 characterizes the agents' behavior. Quite intuitively, agents with low costs of effort participate and comply. Those with large benefits and high effort costs participate but violate. The remaining agents (those with low benefits and high effort costs) abstain from participation.

The next two propositions show how changes in selective incentives affect the absolute number of complying agents ("compliance"), the absolute number of violating agents ("violations"), overall participation (compliance plus violations), and the ratio of the number of compliers over the number of participants ("compliance rate"). The results of both propositions are summarized in Table 7.

Proposition 1. *An increase in the accuracy in identifying violations, q_v, results in an increase in compliance, a decrease in violations and participation, and an increase in the compliance rate. An increase in the accuracy in identifying compliance, q_k, results in an increase in compliance, a decrease in violations, an increase in participation, and an increase in the compliance rate.*

Proof. Proposition 1 can be easily proved by the comparative statics for the three relevant thresholds thresholds b_v, e_k and e_p reported in Table 6, which is trivial to verify. □

Policy variable	b_v	e_k	$e_p(b)$
q_v	Increase	Increase	Constant
q_k	Constant	Increase	Increase
c	Decrease	Increase	Increase
s	Increase	Increase	Decrease

Table 6: Comparative statics

An increase in accuracy has the plausible effect of increasing compliance and reducing violations. This is a well-known result (Becker 1968, Png 1986). However, we also show that the effect on participation depends on whether accuracy increases with respect to identifying violators or compliers. Greater accuracy in identifying compliers, q_k, increases the participation and compliance thresholds but has no effect on the benefit threshold. As a result, some former violators start complying, and some former non-participants do the same, but the decision between non-participating and violating is unaffected. So higher q_k leads to greater participation because some former non-participants start complying and that is the only change made by former non-participants.

In contrast, greater accuracy in identifying violators, q_v, increases both the compliance threshold and the benefit threshold, while keeping the participation threshold unchanged. This means that some former violators start complying, and some former violators stop participating, but the decision between non-participating and complying is unaffected. So higher q_v leads to lower participation because the only change in participation comes from some violators deciding to abstain from participating.

Proposition 2. *An increase in the carrot, c, results in an increase in compliance and participation, while the effect on violations and the compliance rate is ambiguous. An increase in the stick, s, results in a decrease in violations and participation, while the effect on compliance and the compliance rate is ambiguous.*

Proof. The results on compliance and participation follow trivially from the fact that e_k and $e_p(b)$ increase in c, while b_v decreases in c. The effect on violations is ambiguous because the violation region shifts over different types as c increases, which in turn is due to the fact that e_k moves up and b_v moves to the left. Hence, the result depends on the relative frequency of those types in the population of agents, that is, on $f(b,e)$.

Let us now examine carrots. The absolute number of violating agents is given by

$$V \equiv \int_{b_v}^{\infty} \int_{e_k}^{\infty} f(b,e)\, de db$$

Therefore, we have:

$$\begin{aligned}
\frac{\partial V}{\partial c} &= -\frac{\partial b_v}{\partial c}\int_{e_k}^{\infty} f\left(b_v,e\right) de \\
&\quad -\frac{\partial e_k}{\partial c}\int_{b_v}^{\infty} f\left(b,e_k\right) db \\
&= \left(1-q_v\right)\int_{e_k}^{\infty} f\left(b_v,e\right) de \\
&\quad -\left(q_k+q_v-1\right)\int_{b_v}^{\infty} f\left(b,e_k\right) db
\end{aligned}$$

If b and e are independently and uniformly distributed on the unit interval—that is, if $f\left(b,e\right) = 1$—the ambiguity persists and we have that $\frac{\partial V}{\partial c} > 0$ iff

$$\frac{1-e_k}{1-b_v} > \frac{q_k+q_v-1}{1-q_v}$$

Concerning sticks, the results on violations and participation follow trivially from the fact that e_k and b_v increase in v, while $e_p\left(b\right)$ decreases in s. The effect on compliance is ambiguous because the compliance region shifts over different types as s increases because e_k moves up and b_v moves to the right; hence the result depends on $f\left(b,e\right)$. More formally, the absolute number of complying agents is given by

$$K \equiv \int_0^{b_v}\int_0^{e_p(b)} f\left(b,e\right) dedb + \int_{b_v}^{\infty}\int_0^{e_k} f\left(b,e\right) dedb$$

Therefore, we have:

$$\begin{aligned}
\frac{\partial K}{\partial s} &= \frac{\partial b_v}{\partial s}\int_0^{e_p(b_v)} f\left(b_v,e\right) de \\
&\quad +\int_0^{b_v}\frac{\partial e_p(b)}{\partial s} f\left(b,e_p\left(b\right)\right) db \\
&\quad -\frac{\partial b_v}{\partial s}\int_0^{e_k} f\left(b_v,e\right) de \\
&\quad +\frac{\partial e_k}{\partial s}\int_{b_v}^{\infty} f\left(b,e_k\right) db \\
&= -\left(1-q_k\right)\int_0^{b_v} f\left(b,e_p\left(b\right)\right) db \\
&\quad +\left(q_k+q_v-1\right)\int_{b_v}^{\infty} f\left(b,e_k\right) db
\end{aligned}$$

where note that the first and the third line cancel each other out because $e_p\left(b_v\right) = e_k$. If b and e are independently and uniformly distributed on the unit interval, we have that $\frac{\partial K}{\partial s} > 0$ iff

$$\frac{1-b_v}{b_v} > \frac{1-q_k}{q_k+q_v-1}$$

which again shows that the ambiguity persists even if the agents' types are uniformly and independently distributed. □

Carrots and sticks have predictable effects on participation, which increases with carrots and decreases with sticks. However, while carrots unambiguously increase compliance, they may or may not reduce the occurrence of violations. The reason is that, by drawing more agents to the activity, carrots also increase the payoff for those who choose to participate and violate the rule. Similarly, sticks unambiguously reduce violations but may or may not increase the number of complying agents. This is because while sticks induce some violators to comply, sticks also push some compliers away from the regulated activity. The balance of these opposing effects depends on values of the various policy parameters (and in the Appendix we show that these results do not depend on a particular distribution of the agents' types and do not disappear if each point on

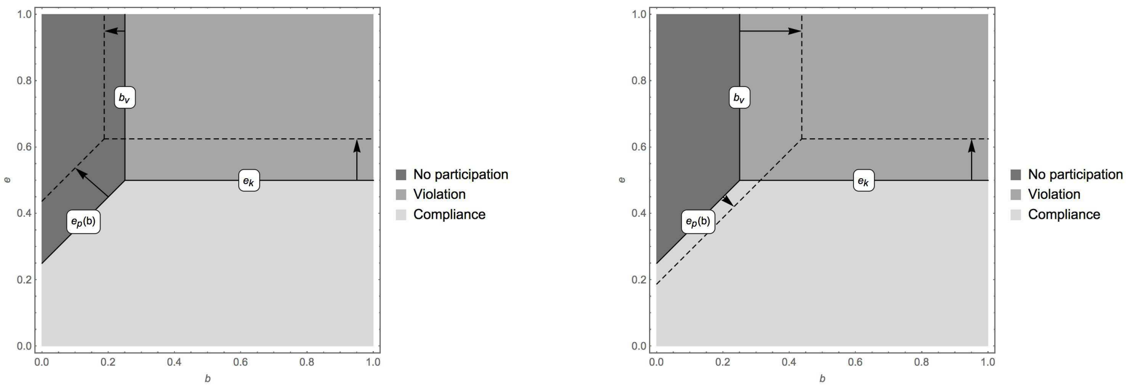

(a) The ambiguous effect of carrots (from $c = \frac{1}{2}$ to $c' = \frac{3}{4}$)

(b) The ambiguous effect of sticks (from $s = \frac{1}{2}$ to $s' = \frac{3}{4}$)

Figure 8: Effects of an increase in the carrot or the stick (simulation parameters: $q_k = q_v = \frac{3}{4}$).

the graph is given equal weight, which is the case, for instance, if e and b are independently and uniformly distributed on the unit interval.)

Figure 8 visualizes these results. An increase in the carrot pushes up both the compliance and the participation thresholds but reduces the benefit threshold b_v. As a result, the compliance area expands, the no-participation area contracts (resulting in more participation), but the effect on the violation area is ambiguous. The increase in the compliance threshold causes some violators to comply. The decrease in b_v, however, enlarges the full participation group, causing some non-participants to become violators.

Similarly, an increase in the stick increases e_k and b_v, but reduces $e_p(b)$. As a result, the violation area contracts, the no-participation area expands (resulting in less participation), while the effect on the compliance area is ambiguous. The increase in e_k induces some violators to comply, while the reduction in $e_p(b)$ leads some compliers to abstain from participation.

Policy variable	Compliance	Violations	Participation	Compliance rate
q_v	Increase	Decrease	Decrease	Increase
q_k	Increase	Decrease	Increase	Increase
c	Increase	Ambiguous	Increase	Ambiguous
s	Ambiguous	Decrease	Decrease	Ambiguous

Table 7: Effects of an increase in the policy variables on compliance, violations, compliance rate, and overall participation

A.1.2 The model with imperfect monitoring ($p < 1$)

In the previous section, we assumed that the principal monitors each agent with certainty. Here, we introduce imperfect monitoring (or auditing) in the model, $p \in (0,1)$, thereby completing the analysis of the general setup of Section (2.1). As a result of imperfect monitoring, agents who are not monitored may–in the general case–receive a carrot with probability $\phi_c \geq 0$, a stick with probability $\phi_s \geq 0$, or neither of the two with the residual probability $1-\phi_c-\phi_s \geq 0$.[23]

To illustrate, if $\phi_c = 1$, then non-monitored agents receive a carrot with certainty. This case corresponds to a subsidy given to agents who engage in the activity and are not found in violation of the rule. Similarly, if $\phi_s = 1$, then non-monitored agents are subject to a stick with certainty. This amounts to taxing the activity unless the agent is found to comply. Finally, if $\phi_c = \phi_s = 0$, then carrots and sticks are applied only to monitored agents. In the general case, non-monitored agents may be taxed or subsidized with some probability.[24]

Action	Payoff
The agent participates and complies	$\Pi_k(b,e) = b + pq_k c - p(1-q_k)s - e + (1-p)(\phi_c c - \phi_s s)$
The agent participates but violates	$\Pi_v(b) = b + p(1-q_v)c - pq_v s + (1-p)(\phi_c c - \phi_s s)$
The agent does not participate	0

Table 8: Agent's payoffs with imperfect monitoring

From the payoffs reported in Table 8, we can calculate the relevant thresholds with imperfect monitoring:

$$\begin{aligned} b_v^i &= p(q_v s - (1-q_v)c) - (1-p)(\phi_c c - \phi_s s) \\ e_k^i &= p(c+s)(q_k + q_v - 1) \\ e_p^i(b) &= b + p(q_k c - (1-q_k)s) + (1-p)(\phi_c c - \phi_s s) \end{aligned}$$

Imperfect monitoring has two effects on the thresholds that we consider in the analysis. First, for all thresholds, imperfect monitoring dilutes the effect of carrots and sticks proportionally to the probability of monitoring p. Second, while the compliance threshold e_k^i does not depend on the treatment of non-monitored agents, the other two thresholds do. The factor $(1-p)(\phi_c c - \phi_s s)$ is the net expected payment that non-monitored agents receive, which is positive if $\phi_c c > \phi_s s$, that is, if non-monitored agents are rewarded in expectation and negative otherwise.

It is easy to verify (see Appendix) that the expanded model including imperfect monitoring generates exactly the same results as the basic model with respect to the selective incentives considered there. In addition, however, when monitoring is imperfect, there is an additional policy variable: the probability of monitoring.

Proposition 3. *Propositions 1 and 2 hold also with imperfect monitoring, $p < 1$. In addition, an increase in the probability of monitoring, p, has ambiguous effects on compliance, violations, participation and the compliance rate.*

Proof. The results follow from the comparative statics reported in Table 9, to be compared with Table 6. Accordingly, we can easily derive Table 10: which proves the proposition. □

As it is easy to verify by inspection, an increase in p results in an increase of e_k^i while having an ambiguous effect on b_v^i and $e_p^i(b)$. In particular, b_v^i increases in p if $(q_v - \phi_s)s > (1-q_v-\phi_c)c$, that is, if the expected stick for violators exceeds their expected carrot, net of all payments to or by non-monitored agents, and decreases in p otherwise. The intuition is that if the expected carrot is large

[23]See Dari-Mattiacci et al. (2009) examining the incentive problems that arise from rewarding or punishing non-monitored agents.

[24]In Section A.2 we examine the classical case of taxes and subsidies that are not conditional to the agent's behavior.

Policy variable	b_v^i	e_k^i	$e_p^i(b)$
q_v	Increase	Increase	Constant
q_k	Constant	Increase	Increase
c	Decrease	Increase	Increase
s	Increase	Increase	Decrease
p	Ambiguous	Increase	Ambiguous

Table 9: Comparative statics

Policy variable	Compliance	Violations	Participation	Compliance rate
q_v	Increase	Decrease	Decrease	Increase
q_k	Increase	Decrease	Increase	Increase
r	Decrease	Increase	Decrease	Decrease
c	Increase	Ambiguous	Increase	Ambiguous
s	Ambiguous	Decrease	Decrease	Ambiguous
p	Ambiguous	Ambiguous	Ambiguous	Ambiguous

Table 10: Effects of an increase in the policy variables on compliance, violations, compliance rate and overall participation with imperfect monitoring

enough, increasing the probability of monitoring induces violators to participate in the activity (due to lower b_v^i). Similarly, the threshold $e_p^i(b)$ increases in p only if $(q_k - \phi_c)\, c > (1 - q_k - \phi_s)\, s$, that is, if compliers face an expected carrot larger than the expected stick. In this case, increasing monitoring induces some violators to start complying. These ambiguities make it impossible to predict the effect of increased monitoring levels on compliance, violations, participation, and relative compliance in the general case.

Let us now turn to a realistic sub-case in which non-monitored agents receive a carrot with certainty, that is, where $\phi_c = 1$ and $\phi_s = 0$. In this case, we have

$$\begin{aligned} b_v^i &= p\,(q_v s + q_v c) - c \\ e_k^i &= p\,(c + s)\,(q_k + q_v - 1) \\ e_p^i(b) &= b + p\,(q_k c - (1 - q_k)\, s) + (1 - p)\, c \end{aligned}$$

It is easy to see that b_v^i now increases in p while $e_p^i(b)$ decreases in p. This resolves some but not all ambiguities. In particular, the effect of an increase in p on absolute and relative compliance remains ambiguous as illustrated in Figure 9.

Corollary 1. *When non-monitored agents receive a carrot then an increase in the probability of monitoring results in a decrease in violations and in participation and has ambiguous effects on compliance and on the compliance rate.*

An increase in monitoring when $\phi_c = 1$ leads to an increase in e_k and b_v, but reduces $e_p(b)$. These are the same shifts that result from an increase in sticks. This similarity is quite intuitive. Greater monitoring amounts to a greater expected cost for violators and mistakenly punished compliers. The resulting consequences and ambiguities are the same as when an increase in the stick raises the same expected cost.

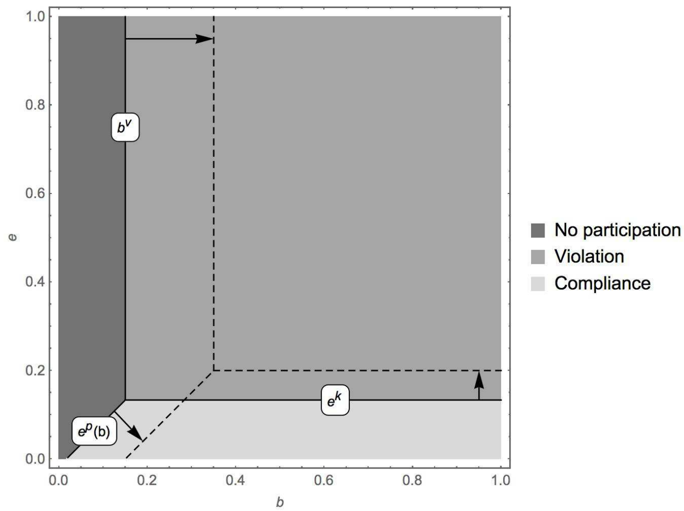

Figure 9: Effects of a change in apprehension rate (from $p = \frac{2}{3}$ to $p^{'} = 1$; other simulation parameters: $q_k = q_v = \frac{3}{5}$, $c = \frac{1}{4}$, $s = \frac{3}{4}$, $\phi_c = 1$, $\phi_s = 0$).

A.2 Taxes and subsidies

Let us now examine whether adding a tax to the model expands the set of outcomes that the principal can reach. This section generalizes the results illustrated in Section (3). The new payoff matrix is as follows:

Action	Payoff
The agent participates and complies	$\Pi_k(b,e) = b + pq_kc - p(1-q_k)s - e + (1-p)(\phi_c c - \phi_s s) + \tau$
The agent participates but violates	$\Pi_v(b) = b + p(1-q_v)c - pq_v s + (1-p)(\phi_c c - \phi_s s) + \tau$
The agent does not participate	0

Table 11: Agent's payoffs with imperfect monitoring and taxes / subsidiesAgent's payoffs

We have now the following thresholds:

$$\begin{aligned} b_v^\tau &= p(q_v s - (1-q_v)c) - (1-p)(\phi_c c - \phi_s s) - \tau \\ e_k^\tau &= p(c+s)(q_k + q_v - 1) \\ e_p^\tau(b) &= b + p(q_k c - (1-q_k)s) + (1-p)(\phi_c c - \phi_s s) + \tau \end{aligned}$$

Clearly the compliance threshold is unaffected by τ, because every participant receives τ whether he complies or not. In contrast, participation will be incentivized by τ, both for compliers and for violators, since τ reduces b_v^τ and increases $e_p^\tau(b)$. The addition of τ, however, does not expand the set of outcomes that the principal can reach. To see why, consider that the principal could replicate the same compliance and participation thresholds as above by simply adjusting the carrot and the stick instead of implementing a tax or a subsidy. In short, the principal can replicate the effects of taxes and subsidies by keeping the sum of the carrot and the stick constant (so as to keep the compliance threshold constant) while altering their ratio $\frac{c}{s}$. That ratio should increase if the activity is subsidized (that is, if τ is positive) and decrease if the activity is taxed (that is, if τ is negative), leading to the following proposition:

Proposition 4. *Complementing carrots and sticks with a tax or a subsidy on the regulated activity does not affect any of the results; the principal can replicate the effect of any tax or subsidy, τ, by setting $c' = c + T$ and $s' = s - T$ where:*

$$T = \frac{\tau}{p + (1-p)(\phi_c + \phi_s)}$$

Proof. Let $c' = c + T$ and $s' = s - T$. It is easy to see that $c' + s' = c + s$ and hence $e_k^i = e_k^\tau$. Note also that the value of T that guarantees that $b_v^i = b_v^\tau$ also necessarily yields $e_p^i(b) = e_p^\tau(b)$. To derive the value of T note that

$$\begin{array}{rcl} b_v^i & = & b_v^\tau \\ & \Longleftrightarrow & \\ p\left(q_v s' - (1-q_v)c'\right) - (1-p)\left(\phi_c c' - \phi_s s'\right) & = & p(q_v s - (1-q_v)c) - (1-p)(\phi_c c - \phi_s s) - \\ & \Longleftrightarrow & \\ p(q_v s - (1-q_v)c) - (1-p)(\phi_c c - \phi_s s) - pT - (1-p)(\phi_c + \phi_s)T & = & p(q_v s - (1-q_v)c) - (1-p)(\phi_c c - \phi_s s) - \\ & \Longleftrightarrow & \\ T[p + (1-p)(\phi_c + \phi_s)] & = & \tau \end{array}$$

which yields the result. ☐

Note that, if monitoring is perfect ($p = 1$) then $T = \tau$. Even if monitoring is imperfect ($p < 1$) we still have $T = \tau$ if $\phi_c + \phi_s = 1$, that is, if non-monitored agents are always either rewarded or punished. The common case where all unmonitored participants receive a carrot is one particular example when this condition is met (see Corollary (1)). In the general case, the adjustment to the ratio of carrots and sticks needs to account for the treatment of non-monitored agents. The intuition is that, if there is a positive probability that a non-monitored agent is subject neither to a carrot nor to a stick, that agent fails to be compensated for the effect of the tax or subsidy and hence the correction on the carrot and the stick needs to take that eventuality into account.

A.3 Agents' perceptions of accuracy

Here we consider an alternative model where all agents have a benefit of participation $b = 0$. Agent's, however, vary in their perception of the accuracy of enforcement by the principal so that the three-dimensional agent type is (e, q_k, q_v), where we assume that e is distributed according to $f(e)$ and is independent of (q_k, q_v) which are distributed according to $g(q_k, q_v)$ and are correct on average, that is, the mean of the distribution is equal to the true levels of accuracy $\left(q_k^g, q_v^g\right)$. Additionally, we assume that a change in q_k^g or q_v^g shifts the distribution according to the single-crossing property. We will show that all the results obtained in the basic setup are confirmed in this alternative setup.

The agents' payoffs are:

As in the basic model, we can define three thresholds:

$$\begin{array}{rcl} \hat{q}_v & \equiv & \frac{c}{c+s} \\ e_k(q_k, q_v) & \equiv & (c+s)(q_k + q_v - 1) \\ e_p(q_k) & \equiv & q_k c - (1-q_k)s \end{array}$$

Action	Payoff
The agent participates and complies	$\Pi_k(e, q_k) = q_k c - (1 - q_k) s - e$
The agent participates but violates	$\Pi_v(q_v) = (1 - q_v) c - q_v s$
The agent does not participate	0

Table 12: Agent's payoffs

The threshold $\hat{q}_v$ is the level of an agent's perception of the enforcement accuracy, q_v, that makes the agent indifferent between violation and abstention, that is, such that $\Pi_v(q_v) = 0$. This threshold is analogous to the threshold b_v in the basic model and partitions the agent's population into two groups: the willing participants, with $q_v < \hat{q}_v$—whose compliance decision is controlled by the compliance threshold $e_k(q_k, q_v)$—and the reluctant participants, with $q_v \geq \hat{q}_v$—whose compliance decision is controlled by the participation threshold $e_p(q_k)$.

If $c > s$, we have $\frac{1}{2} < \hat{q}_v < 1$. In this case, the threshold $\hat{q}_v$ lies in the interval of admissible values for q_v. This implies that, stochastically, some agents are in the full participation group and others are in the partial participation group. If instead $c \leq s$, then $\hat{q}_v \leq \frac{1}{2}$ and hence all agents are in the group with $q_v > \hat{q}_v$, that is, all agents are in the partial participation group and either participate and comply or stay away from the regulated activity. Figure 10 illustrates how agents of different types behave when $c > s$.

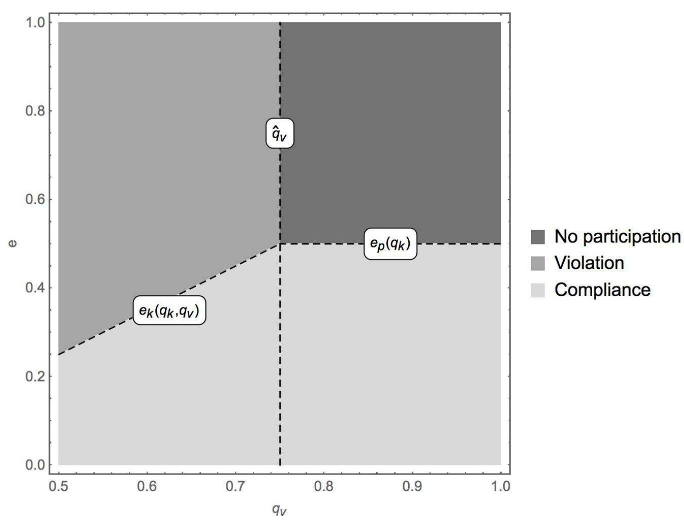

Figure 10: Agent's behavior (simulation parameters: $q_k = \frac{3}{4}$, $c = \frac{3}{4}$, $s = \frac{1}{4}$).

We are interested in establishing the expected levels of compliance and participation in the population. We start by summarizing the comparative statics of the thresholds of interest in Table 13. Recall that here, while the first three variable are policy variables, q_k is an agent's characteristic and varies across agents.

Accordingly, we can derive the results in Table 14, which are in accordance with those obtained in the basic model.

Policy variables	$\hat{q}_v$	$e_k(q_k, q_v)$	$e_p(q_k)$
c	Increase	Increase	Increase
s	Decrease	Increase	Decrease
q_k	Constant	Increase	Increase

Table 13: Comparative statics

Policy variable	Compliance	Violations	Participation	Compliance rate
q_v^g	Increase	Decrease	Decrease	Increase
q_k^g	Increase	Decrease	Increase	Increase
c	Increase	Ambiguous	Increase	Ambiguous
s	Ambiguous	Decrease	Decrease	Ambiguous

Table 14: Effects of an increase in the policy variables on compliance, violations and overall participation with imperfect monitoring

While the last three lines in Table 14 can be trivially derived from Table 13, the first two lines require a more involved derivation. The result on the effect of q_k^g follows from the comparative statics in Table 13 and the single-crossing property: an increase in q_k^g results in a larger probability mass on higher values on q_k, which in turn correspond to higher values of the compliance and participation thresholds, yielding the result. Finally, the result on the effect of q_v^g follows from the single-crossing property: an increase in q_v^g results in a larger probability mass on higher values on q_v; by inspecting Figure 10 it is easy to see that as we move to higher levels of q_v compliance weakly increases, violations decrease (initially in a continuous fashion and then discontinuously to zero) and participation increases (in a discontinuous fashion), which yield the result.