# *Regulation with Interested Experts*

Jason Scott Johnston*

Draft of May, 2012

ISNIE Discussion Draft (Incomplete: please do not distribute beyond conference participants)

## Abstract

This article constructs a sequential game theoretic model in which regulation is based on expert reports regarding a state of the world --- such as the social benefits from regulation – that is known to the expert but not observed by a decision maker. Experts are assumed to have policy preferences over regulation. When the expert's preference is known but experts incur a positive cost to report to the regulator (or court), it is shown that there is a unique sequential equilibrium in which a pro-regulation expert falsely reports that regulatory benefits are high with a higher probability, the larger are ex ante regulatory benefits; in cases where a net benefit maximizing regulator is close to being ex ante indifferent between regulating or not, the pro-regulation expert is most truthful (in the sense of reporting honestly with a high probability). Two remedies for expert bias are considered: imperfect auditing by the regulatory decision-maker, and expert competition. It is shown that imperfect auditing increases the reliability of expert advice, and the more imperfect the audit, the more reliable is the expert advice. Expert competition between experts with conflicting preferences likewise always increases the information available to the decision-maker, although at least one expert will generally have to be compensated to provide advice against her interest. The result is generalized to the case where expert preferences are known only to the expert and where the expert must incur a positive cost to become perfectly informed with some positive probability. In this case, expert effort to become informed is maximized by the expert who is just indifferent between regulating and not ex ante. For this reason, the optimal expert from the point of view of the decision maker generally has preferences that lie between those of the decision maker and those which maximize expert effort.

## *I. Introduction*

Virtually all modern regulators rely upon expert advice in deciding whether and how to regulate. Environmental regulations proposed by the U.S. Environmental Protection Agency, for example, are most commonly justified by the benefits they will bring in terms of a reduction in premature deaths (typically due to cancer) caused by exposure to various pollutants. The climate change regulations proposed by that same agency are more broadly justified by the many harms that the Intergovernmental Panel for Climate Change (IPCC) Assessment Reports says will be made less likely if human greenhouse gases are quickly reduced. Outside the environmental regulatory realm, U.S.

* Henry L. and Grace Doherty Charitable Foundation Professor of Law, Nicholas E. Chimicles Research Professor in Business Law and Regulation, University of Virginia School of Law. Email: jjohnston@virginia.edu.

monetary policy is based on what economists say about the likely consequences of those policies for the rate of inflation and level of unemployment; financial regulation about what economists and financial advisers of one stripe or another predict will be the consequences of alternative regulations for bank solvency, stability and (perhaps) lending practices.

Two beliefs about expert based regulation have become institutionalized: the first is that experts are not interested in regulatory decisions, but are disinterested purists who truthfully tell regulators, or decision makers more generally, what they know about the relevant state of the world (e.g. health benefits from EPA regulations). The second and corollary belief is that decision makers who are not themselves experts should be highly deferential to expert advice and simply follow what the experts recommend.

The position favoring extreme deference to expert advice remains an important part of U.S. administrative law. The U.S. Supreme Court has said that courts should be at their "most deferential" when reviewing a regulatory agency's assessment of the science supporting regulation.[1] Although lower federal courts have sometimes been much less deferential to agency science than this standard requires,[2] the degree of deference they accord is uncertain and sometimes seems to approach the extreme deference standard.[3] According to the former President of the Sierra Club Legal Defense Fund, such an attitude has been interpreted by many federal judges to mean that the law "requires that conflicts among experts always be resolved in favor of the government. Always means that even if you have one hundred experts holding an opinion in your favor while only one expert supports the government, you still lose."[4]

The justification for extreme deference to expert assessments of regulatory science has been well stated by the late activist climate scientist Stephen Schneider and his political scientist colleague Paul Edwards. As they describe it, in the climate science field, disagreement among scientists, not just on "details," but even sometimes "major points," is an "unavoidable element of a still-inexact science."[5] But the assessment of what are "best range of probability and range of consequences estimates" for climate change are not to be made by either the general public or lay politicians, for "they cannot be expected to determine for themselves how to weigh these conflicting [scientific] opinions."[6] Instead, the ultimate assessment of what are the "best" scientific estimates are to be made by the "responsible scientific community."[7]

---

[1] Baltimore Gas & Electric Co v. NRDC, 462 U.S. 87, 103 (1983).

[2] *Ecology Ctr., Inc. v. Austin* 430 F.3d 1057 (2005) (overruled by en banc in *Lands Council v. McNair,* 629 F.3d 1070 (9th Cir. 2010).

[3] See, for example, *Lands Council v. McNair*, 629 F.3d 1070 (9th Cir. 2010).

[4] Vic Sher, Breaking out of the Box: Toxic Risk, Government Actions, and Constitutional Rights, 13 J. Envtl. L. & Litig. 145, 148-149 (1998).

[5] Paul N. Edwards and Stephen H. Schneider, Self-Governance and Peer-Review in Science for Policy, in Changing the Atmosphere 219, 244 (Clark A. Miller and Paul N. Edwards, eds. 2001).

[6] Id. at 243.

[7] Id. at 244.

The problem with this argument for deference is that it rests on a belief in the existence of a "responsible scientific community" of disinterested experts. There is, however, virtually no reason to think that the belief in disinterested experts is true in the world of regulatory scientific advice. In many regulatory areas, experts have very strong preferences over alternative regulatory outcomes. There are many examples of such strong preferences. In criticizing the U.S. Endangered Species Act's requirement that decisions whether or not to list species as endangered or threatened is to be made "solely on the basis" of the best available scientific information,[8] Holly Doremus argues that "regulated interests may be able to influence regulatory evaluation through submission of flawed or misleading scientific information…anytime the regulatory system is dependent upon information that is peculiarly within the control of those with a strong financial interest in the outcome of the regulatory process there very real opportunities for scientific review to be hijacked or skewed."[9] Notably, she recognizes that this kind of slide from "skeptical evaluation" to "advocacy" can be triggered not only by employment by a regulated entity, but also by "other sorts of strong policy preferences."[10] Such policy preferences are illustrated by the field of conservation biology, a field whose practitioners provide crucial expert advice to natural resource regulators who are responsible for determining how to manage public lands so as to further the legislative goal of protecting biodiversity. According to one of the founders of conservation biology, "the entire field rests on the value assumption that biodiversity is good and ought to be conserved. Human actions that protect and restore biodiversity are good; those that destroy or degrade biodiversity are bad…," and conservation biologists "have an ethical obligation to make a powerful case for the conservation of biodiversity to everyone, everywhere."[11] A survey of the attendees of the 2006 meeting of the Society for Conservation Biology found that "70% felt that the voice of the society – the journal *Conservation Biology* – should advocate certain policies."[12]

In the area of climate change regulation, many individuals delivering expert advice have also had obvious and strong policy preferences in favor of regulations requiring reductions in greenhouse gas emissions. Examples include: Anthony Socci, who for many years was the American Meteorological Society's climate science spokesperson, was a former staffer to Al Gore; the late John Firo, who as administrative director of NCAR was a frequent "expert" commentator on global warming, also served for five years as Board Chairman of the environmental advocacy group Environmental Defense; the UK Meteorological Office's chairman, Robert Napier, was previously the Chief Executive for the World Widlife Fund UK; John Holdren, who before his current service in the White House Science Advisor served as a Clinton-Gore administration spokesperson on global warming; Michael Oppenheimer, a lead author on IPCC Reports, was once the Barbara Streisand scientist at Environmental Defense.[13]

---

[8] 16 U.S.C. §1533(b)(1)(A) (2000).
[9] Doremus, 86 Texas L. Rev. at 1617.
[10] Doremus, 86 Texas. L. Rev. at 1617.
[11] Reed Noss, Values are a Good Thing in conservation biology, 21 Conserv. Bio. 18 (2007)
[12] Emma Marris, Should Conservation Biologists Push Policies?, 442 Nature 7098 (2006).
[13] See Richard S. Lindzen, Climate Science: Is it Currently Designed to Answer Questions? 5-8(2008).

Acknowledging that experts often do indeed have such preferences over policy outcomes is not to deny expertise. But it does raise the question of whether the fact that experts have policy preferences ought in some way to influence the degree of deference given by decision makers to expert advice. To gain some insight into this question, this article constructs a simple game theoretic model in which regulation may be based on expert reports regarding a state of the world --- such as the social benefits from regulation – that is known to the expert but not observed by a decision maker. The article explores the benchmark case where the decision maker wishes to regulate only if regulatory benefits are larger than the (commonly known) cost of complying with the regulation, but the experts reporting to her have extreme policy preferences, either favoring or opposing regulation no matter what the actual social benefits. At first, I assume that the decision maker cannot even imperfectly check on the truth or falsity of the expert's report. Unlike many (but not all) economic models of expert advice, I assume that the expert incurs a cost to prepare and communicate her report to the decision maker, and that this cost is higher when the expert's report is false than when it is true. That experts do incur a cost when they attempt to assess and report on the current state of scientific knowledge about things such as the harm from climate change or the risk from chemical exposure is obvious from all accounts of the expert reporting process. My assumption that costs are higher when the expert falsely reports may be rationalized in a number of ways: as representing the expert's own psychic pain incurred when she violates her professional standards and communicates false information, or as the differential cost of preparing a biased versus accurate report.

I first solve for Perfect Bayesian Equilibria (PBE) in the expert reporting game when the decision maker cannot audit the expert. BPE are strategies adopted by the expert and decision maker such that each party's strategy is payoff maximizing given her beliefs and the strategy of the other player, and given that beliefs are consistent with Bayes' Rule wherever possible (viz. on the equilibrium path).[14]

In the expert reporting game analyzed below, there are some PBE with no expert report. These occur when the choice that is optimal for the decision maker without any advice from an informed expert is the one that the expert prefers. But when the decision maker's uninformed decision is one that the expert would not prefer, the expert will report to the decision maker. Such reports are always truthful if the true state of regulatory benefits is the value that leads to the outcome preferred by the expert (high benefits for the pro-regulation expert, low benefits for the anti-regulation expert). But when the true state of the world is the one that the expert does not prefer, then the expert sometimes – that is, with a probability between but not equal to 0 and 1 – reports falsely. This possibility of an informative report by the expert causes the expert to regulate with a positive probability when she would not regulate if uninformed, and to refrain from regulating with positive probability when she would regulate if uninformed. In equilibrium, the probability that a pro-regulation expert falsely reports that regulatory benefits are high when they are actually low depends upon the ex ante case for regulation. If costs and benefits given only prior probabilities are such that regulation is almost

[14] This succinct definition is given by Joel Watson, Strategy: An Introduction to Game Theory 342 (2d ed. 2008).

optimal for the ideal regulator, then even a very biased expert report –with a high probability of a false report – can induce regulation. However, when the ex ante case tilts strongly against regulation, the pro-regulation expert can induce regulation only if her report is very reliable, in the sense of a high probability of truthfully reporting that benefits are low so that regulation does not occur. As for the probability that regulation occurs, this is determined, intuitively enough, by the costs and benefits to the expert of inducing regulation with a false report. For an expert who perceives a very big cost from false reporting, regulation has to occur with a high probability for the expert to be willing to sometimes bear that cost.

The inverse conditions apply for an expert who has anti-regulatory preferences and internalizes only the cost but none of the benefit of regulation. Regardless of the direction of expert bias, however, the decision maker may encounter a quite high probability of a false expert report in equilibrium (and still be better off using this report than not). This result generalizes to the situation where there is a continuum of expert types with preferences that vary from strongly pro- to strongly anti-regulation and where the expert must decide whether to incur a positive effort cost to become informed about the state of the world. It is shown that effort is maximized by an unbiased expert, one who shares the regulator's preference. The result on the reliability of expert advice is generalized to this case, and it is shown that for very low prior expected regulatory benefits, an expert recommendation to regulate is very likely informed and therefore very reliable. By contrast, when regulatory benefits are ex ante expected to be high, there is a high likelihood that an expert recommendation to regulate is uninformed, and so expert advice is less reliable.

These results raise the obvious question of whether the decision maker should be deferential to the expert report, or instead attempt to make her own assessment of the state of regulatory benefits, and regulate (or not) only if her own assessment agrees with the expert's report. It turns out that provided that the decision maker's own signal is informative, albeit imperfectly, of the state of the world, then she should indeed condition her decision on agreement between her signal and the expert's report. However, quite non-intuitively, in the equilibrium that obtains when the decision maker is not deferential to the expert, and audits the expert with her own signal, the expert becomes more honest, the more imperfect is the decision-maker's signal. The reason has entirely to do with the structure of PBE in this game: a pro-regulation expert who wishes to induce a positive probability of regulation realizes that the more imperfect is the decision maker's signal, the more reliable must be her own report for the decision maker to regulate. Essentially, the more knowledgeable is the decision maker, the less she depends on the expert report, and the more biased becomes that report. Consequently, when the decision maker is not deferential to the expert, but uses her own beliefs as well, it is the decision maker who very often errs in determining the actual state of regulatory benefits – the classic uninformed layperson -- who induces a very high probability of a truthful expert.

An alternative response to expert bias is to pit one type of biased expert against another. This version of the paper concludes by analyzing this institutional setup. It is shown that the decision maker always achieves a higher expected return when she employs two experts with conflicting preferences than when she uses just one, and that each type of expert is induced to be more honest by such expert competition. However, the existence of a PBE in mixed strategies may require that at least one expert be paid when she reports against her interest.

This initial model assumes an expert whose preferences are known and who has perfect knowledge of the true state of regulatory benefits, and who incurs a cost only to prepare and communicate her report to the decision maker. In the second part of the paper, the analysis is extended to the case where expert preferences vary, and where the expert choose how much effort to expend to – possibly – become informed as to the true state of regulatory benefits. Expert types are distinguished by the weight they attach to regulatory compliance costs, with some experts overweighting and others underweighting this cost, while the decision maker is assumed to be a net benefit maximizer who attaches a weight of 1 to both costs and benefits. Using a framework developed in Dur and Swank,[15] it is shown that effort to become informed is maximized by an expert type who is ex ante indifferent between regulating and not, given the ex ante expected regulatory benefits and her weighting of (known) compliance costs. As effort increases the probability that the expert perfectly learns regulatory benefits, it is the expert who views regulation as a close case ex ante who provides the most reliable advice. This is true when expert advice is always followed by the decision maker. However, if the expert's preferences diverge too far from those of the decision maker, then the decision maker will not follow the expert's advice, and in equilibrium such an expert will not incur positive cost to become informed. Hence, there is a tradeoff between expert effort and expert preferences: the optimal and equilibrium expert is one with preferences somewhere between those of the decision maker (attaching weight of 1 to compliance costs) and those which generate expert ex ante indifference and therefore maximize expert effort. When the ex ante case for regulation is close, in that expected benefits are close to compliance costs, the optimal expert is one with preferences vary close to those of the decision maker, and her advice is highly informed and reliable. When the ex ante case for regulation is either very weak or very strong, the decision maker faces an inevitable tradeoff and settles for an expert who is both somewhat unreliable – in that her preferences diverge from those of the decision maker – and relatively likely to have an uninformed opinion about regulatory benefits. This result is somewhat analogous to that obtained for the case when an expert of known pro-regulatory bent maximizes the probability with which she recommends regulation: there too, expert advice is most reliable when the ex ante case for regulation from the decision maker's point of view is close and least reliable when the ex ante case for regulation is very strong.

## *I. Equilibria in Expert-Based Regulation when Experts are Biased*

Let the benefit from the regulation be given by Θ, with Θ initially assuming one of two values, $\Theta_h$ and $\Theta_l$, with $\Theta_h > \Theta_l > 0$. Let the decision maker's perceived cost of regulation be given by $k > 0$, and assume that $\Theta_h > k > \Theta_l$. The cost k may be thought of

[15] Robert Durr and Otto H. Swank, Producing and Manipulating Information, 115 *Econ. J.* 185-199 (2005).

as the compliance cost of regulation. Hence, regulation is optimal from the point of view of the decision maker if and only if $\Theta = \Theta_h$. The prior probability that $\Theta = \Theta_h$ is given by $p$. Let $f(x|y)$ denote the conditional likelihood function, that is, the likelihood that the expert reports $x$, with $x = H, L$, given that the true state of the world observed by the expert is given by $y = h,l$. Accordingly $f(x,y)$ is the joint likelihood function of report x and true state of the world y induced by an expert reporting strategy $f(x|y)$. Let $g(y|x)$ denote the posterior probability that the true state of the world is $y$, given that the expert has reported a state $x$.

The structure of the game is as follows. Nature moves first and chooses the state of the world – low or high regulatory benefits. The expert has perfect information about $\Theta$ (this may be thought of as the state of existing knowledge about regulatory benefits). The expert then either submits a report on the state of regulatory benefits to the regulator (participates in the expert reporting game) or does not. Her cost of preparing and communicating the report depends upon whether or not she reports truthfully. An expert who does not report truthfully – who tells the regulator that benefits are "high" when they are actually low, or vice versa -- has a higher reporting cost than one who reports truthfully. This may be justified either by thinking of the expert's own psychic harm from violating her principles and lying about what she knows, or as the additional cost incurred to prepare a report that manages to look truthful but which in fact is false in its conclusion. For simplicity normalize the cost of truthful reporting to 0, and let $c$ be the cost to the expert of submitting an untruthful report. This conclusion is a statement of the form that $\Theta = \Theta_H$ or, alternatively, that $\Theta = \Theta_L$. Assume that $c < \Theta_L$. One may think of this as the expert reporting either that benefits are "high" or that they are "low." After getting the expert report, and with perfect information regarding expert preferences, the decision maker makes the dichotomous decision to either regulate or does not. If the expert does not participate in the export reporting game, her payoff is 0, which is also the decision maker's status quo payoff from no regulation.

A. Unbiased Experts

As a benchmark, suppose that it is common knowledge that the decision maker will regulate if and only if the expected benefits of regulation exceed the decision maker's perceived cost k. It is also common knowledge that the expert has no concern with whether or not regulation happens, but feels herself bound by professional norms to simply honestly report whatever she observes about regulatory benefits. In this case, it is trivial to see that the decision maker will regulate if and only if the expert reports "high." Given that the decision maker regulates with probability 1 after being given a report communicating that benefits are high, the trusts the expert to be honest, but as the expert has no payoff from dishonesty, the decision maker's trusting strategy is optimal.

B. . Pro-regulation experts

Suppose instead that the expert has different preferences, and perceives a payoff equal simply to the gross benefit from regulation, $\Theta$. In this case, there is a divergence in interest between the decision-maker , who wishes to regulate if and only if $\Theta = \Theta_H$, and the expert, who wishes to regulate regardless of what she observes, as even $\Theta_L > 0$.

Assuming that preferences of both the decision maker and the expert are common knowledge to both, the primary behavioral question is whether the expert can succeed in using her report to generate regulation, even when the decision maker would prefer not to regulate. There are two cases to consider: the first, where the decision maker has no other source of information, other than the expert, about the true state of the world. In the second case, the expert can audit the expert, meaning that she can get her own signal of the true state of the world, after the expert reports, and base her decision on both the expert's report and the value of her independent signal. In this section, I consider the benchmark case where the expert is the decision maker's only source of information as to the state of the world.

For a sufficiently high prior probability $p$ that regulatory benefits equal $\Theta_H$, and sufficiently large benefits $\Theta_H$, there exists a Perfect Bayesian Equilibrium (PBE) in which the decision maker will regulate on the basis of her prior knowledge alone and the pro-regulation expert does not participate. To show this is true, it is enough to see that when regulation is ex ante optimal, we have:

$p\Theta_H + (1 - p)\Theta_L > k$, or if,

$$p > (k - \Theta_L)/(\Theta_H - \Theta_L). \tag{1}$$

Hence as claimed, for a sufficiently high prior probability that regulatory benefits are large, and sufficiently big regulatory benefits in this state of the world, the decision maker goes ahead and regulates without any expert advice. A pro-regulation expert would then never incur positive cost to induce regulation, since it occurs without her participation. But the expert cannot avoid this cost by always reporting the true state, because then regulation would not occur if she reported that benefits were low. Hence when the ex ante decision is top regulate, the decision favored by the pro-regulation expert, such an expert will not participate.

Suppose to the contrary that inequality (1) does not hold. When this is true, the prior probability and large regulatory benefit are not sufficiently high for the decision maker to proceed with regulation based only on her prior knowledge about benefits. When (1) does not hold, the expert may have an incentive to incur the positive cost of untruthful reporting, if such a report can induce a positive probability of regulation. But the expert must provide some information to alter the decision. That is, the expert's objective is to maximize the probability of regulation, since she always gains from regulation. However, if the regulator always reports that benefits are high, then her report is uninformative and regulation does not occur. As the expert gains the most when benefits are actually high, she always truthfully reports high benefits when this is indeed true. But when regulatory benefits are low, the expert must report "high" with some probability less than one for regulation to occur.

In other words, the equilibrium in this game is one in which the expert chooses a mixed strategy of sometimes reporting "high" when the true state is low, and the regulator always regulates. The key probability in this equilibrium is $f(H|l)$, the probability that the expert reports "high" when the true state is low. To solve for this probability, begin with the decision maker. If as in the conjectured equilibrium, the expert pursues a strategy in which $f(H|l) > 0$ and so $f(L|l) < 1$, then the decision maker's expected payoff is given by:

$$g(h|H)\Theta_h + g(l|H)\Theta_l - k. \quad (2)$$

Recall that the decision maker's status quo payoff is *0*. Given this, the decision maker is indifferent between regulating and not whenever her expected payout from regulating is equal to 0. Assuming that the decision maker optimally updates her beliefs according to Bayes' Theorem, and that the expert always reports "high" (or "*H*") when the true state is indeed high (that is, *h*), we have that:

$$g(h|H) = p/[p + (1\text{-}p)f(H|l)] \text{ and}$$

$$g(l|H) = f(H|l)(1\text{-}p)/[p + (1\text{-}p)f(H|l)]. \quad (3)$$

Substituting (4) into (2) and setting (2) equal to 0, we have:

$$\frac{p}{p+(1-p)f(H\mid l)}\theta_h + \frac{f(H\mid l)(1-p)}{p+(1-p)f(H\mid l)}\theta_l = k\,,$$

which simplifies to:

$$f^*(H\mid l) = \frac{p(\theta_h - k)}{(1-p)(k-\theta_l)}\,. \quad (4)$$

Equation (4) derives the probability with which the expert falsely reports that benefits are high that just makes the decision maker indifferent between regulating and not, given an expert report of high benefits. Observe also that the value for $f(H|l)$ given by (4) is the highest such probability consistent with a positive probability of regulation; it is therefore the unique payoff-maximizing BPE for the expert.

As can be seen from (4), the equilibrium probability that the expert falsely reports high benefits, $f^*(H|l)$, increases:

i) the higher is the decision maker's prior probability of high benefits,

ii) the higher is the net benefit from regulating when benefits are indeed high, and

iii) the lower is the net regulatory cost of erroneously regulating when benefits are low.

As can also be seen from (4), existence of such an equilibrium mixed strategy requires that the ex ante or prior expected net benefits from regulating given high benefits are less than the ex ante expected net costs of regulating given low benefits. (For as shown earlier, if this is not true, then we have the uninformative equilibrium).

The intuition behind equation (4) is quite direct. The closer is the ex ante case for regulation, even extremely biased expert advice can successfully induce regulation. To see this most dramatically, suppose that the decision maker is uninformed ex ante in the sense that $p = (1 - p) = 1/2$. In this situation, we see from (4) that as $(\Theta_h - k)$ approaches

*($k$ - $\Theta_l$)*, so that regulation was almost desirable even ex ante, the expert becomes almost completely biased in favor of recommending regulation (more precisely, $f(H|l)$ goes to *1*). Somewhat non-intuitively, the expert becomes more honest and less biased, the weaker is the ex ante case for regulation, the outcome that the decision maker prefers.

Consider next the probability of regulation. If the pro-regulation expert is to play the mixed reporting strategy $\{f^*(L|l), f^*(H|l)\}$ identified above, then it must be that the probability of regulation is such that the expert is indeed indifferent between the two reports, given that the true state is low. Let the probability of regulating given that the expert reports "high" be given by $r^*$. This probability must be such that the expert is just indifferent between reporting "low" versus "high" given that the true state is low, which requires that:

$$0 = [r^*(\theta_l - c) + (1 - r^*)(-c)] \text{, or,}$$

$$r^* = \frac{c}{\theta_l}. \tag{5}$$

Equation (5) says that the lower is the expert's cost of an untruthful report, and the bigger is the gross benefit to the expert from making such a report, the lower is the required probability of regulation. An expert who can expect a large net gain from reporting that benefits are "high" when they are in fact low will be indifferent between that dishonest report and an honest report generating a 0 net return even when there is a very low probability of receiving the net gain from regulation achieved by a dishonest report.

Observe also that if the decision maker would not regulate based on her prior probabilities alone (that is, (1) does not hold), then the pro-regulation is always better off participating in the expert report game than not participating. This is because regardless of what the expert reports – "low" when this is really the case, or "high," triggering regulation with probability $r^*$ -- the expert gets an expected payout of 0, the same that she gets when she does not participate.

C. Anti-regulation Experts

Consider now the polar case of an expert who perceives *0* benefit from regulation and cares only about the compliance cost $k$. More precisely, such an experts gets a gross payoff of 0 from no regulation and $-k$ from regulation. For such an expert regulation is never optimal, and her payoff from participating in the expert reporting game is the compliance cost $k$ that she avoids by participating. Hence such an expert will participate if and only if (1) does not hold, so that the decision maker would indeed regulate based solely on a priori knowledge. As $k$ is the expert's potential gain from forestalling regulation, she will participate only if $c < k$, which means that her gain from participating (avoiding the cost $k$) is bigger than her cost of participating and reporting dishonestly, $c$.

The anti-regulation expert of course always reports truthfully when benefits are low, because in that state her preferences are aligned with the decision maker's, who does not regulate when she knows that benefits are low. When regulatory benefits are known by the expert to be high, however, then she has an interest in deterring regulation. The equilibrium probability that the anti-regulation expert falsely reports that regulatory benefits are low when they are actually high $f_0*(L|h)$ is the highest such probability consistent with the expert preserving sufficient credibility that her report moves the decision maker back to a state of indifference between regulation and no regulation. This is given by:

$$f_0 *(L \mid h) = \frac{(1-p)}{p} \frac{k - \theta_l}{\theta_h - k}. \tag{6}$$

From (6), we see that the conditions determining the honesty of the anti-regulation expert are simply the inverse of those determining the honesty of a pro-regulation expert. Once again, the anti-regulation expert is most reliable (in the sense of low $f_0*(L|h)$) when the ex ante case for regulation from the ideal decision maker's point of view is strongest. The reason is that when the ex ante case for regulation is strong, then an anti-regulation expert can forestall regulation only if her report on regulatory benefits is known to almost always be honest.

For the expert to play the mixed strategy identified in (6), we again require a probability of regulation, denoted by $r_0^*$, that makes the expert just indifferent between a truthful and honest report – now, for the anti-regulation expert – given that the true state is high benefits. Such an $r_0^*$ solves:

$-k = r_0^*(-k-c) + (1-r_0^*)(-c)$, which yields,

$$r_0^* = \frac{k-c}{c}. \tag{7}$$

Since it must be that $r_0^* \in (0,1)$, such an $r_0^*$ exists iff $k < 2c$. More generally, we see from (7) that if it is very costly to the expert to provide an untruthful report, then an anti-regulation expert can be useful to the decision maker – in the sense of sometimes providing a truthful report that is against her preferences – even if there is a relatively low probability of the regulatory outcome. Conversely, if the regulatory compliance cost internalized by the anti-regulation expert is very high, then she will only sometimes provide a truthful report against her interests only if the probability of regulation is very high.

Observe finally that the anti-regulation expert will always participate in the expert reporting game if regulation would occur without her report. This is because regardless of whether she (honestly) reports that the state is high or reports that the state is low, she gets the same expected payout from participating in the reporting game as she would get is she did not participate, $-k$.

*II. Responses to Expert Bias*

As neither type of biased expert is always honest in her report as to the state of the world, the ideal decision maker might well ask whether she could do gain more information from an alternative institutional structure. In this part, I consider two such alternatives: getting her own signal of the state of regulatory benefits, and regulating only if both her signal and the expert report agree that the state is "high;" and, alternatively, getting advice from two experts with opposing preferences.

A. Costless but Imperfect Auditing

To model the possibility of a more active, auditing decision maker, I assume that after the expert reports, the decision maker costlessly observes a signal, *s*, that is correlated with the true state of the world. The signal takes values *h* or *l*, just as does the true state. Let $s_i^j$ = probability that the signal takes value *i,* given that the true state of the world is *j*, for *i,j = h,l.* I shall assume that the signal is informative meaning that for *i,j = h, l,* we have that $s_i^i > s_i^j$. As the results are qualitatively analogous regardless of the direction of expert bias, I consider only the case where the expert has a pro-regulation bias.

The first result is this: provided that the signal is informative in the sense just defined, the decision maker can always achieve a higher equilibrium payoff in the case where she audits and gets to observe the signal *s* than she can in the no audit case. The decision maker achieves this higher payoff by regulating if and only if both the expert says *"h"* and her own signal *s=h.*

To show this, begin by noting that under this presumed regulatory strategy, the probability of regulation is given by the joint probability *Pr(s=h, expert reports h)*. But as there are two states of the world only, we have that *Pr(s=h, expert reports h) = Pr(s=h, expert reports h, true state is h) + Pr(s=h, expert reports h, true state is l)*. But then by the rules on conditional probability and Bayes' Theorem, we have (eventually substituting the notation defined above), that:

$$Pr(s=h, \text{expert reports } H, \text{true state is } h) + Pr(s=h, \text{expert reports } H, \text{true state is } l)$$

$$= Pr(s=h|h)f(H|h)Pr(h) + Pr(s=h|l)f(H|l)Pr(l)$$

$$= s_h^h f(H|h)p + s_h^l f(H|l)(1-p),$$

$$= ps_h^h + (1-p)s_h^l f(H|l),$$

where we used the fact under any optimal, pro-regulation expert strategy, $f(H|h)=1$.

Next observe that under the presumed regulatory strategy, the decision maker's expected payoff is given by:

$g(h|s=h, H)\Theta_h \; + \; g(l|s=h,H)\Theta_l.$

$= g(h|s=h, H)(\Theta_h - \Theta_l) \; + \Theta$ (8)

We have, however, that:

$g(h|s=h, H) = \; Pr(h, s=h, \text{expert reports } H)/Pr(s=h, \text{expert reports } H)$

$= ps_h^h \,/\, [ps_h^h \; + (1-p)s_h^l f(H|l)].$ (9)

Next compare the regulatory expected payoff when the decision maker has available the signal s, which she uses in the presumed strategy – regulating only if the expert's report of h is confirmed by her own signal's value – with the regulatory payoff with no such auditing signal available. As we can see by comparing (8) with (4), the decision maker's expected payoff is higher when she audits and uses the signal s to check the expert than when she does not audit and gain such an independent signal provided that:

g(h|s=h,H) > g(h|H), which using (9) and (3) becomes:

$$\frac{ps_h^h}{ps_h^h + (1-p)s_h^l f(h \mid l)} > \frac{p}{p + (1-p) f(h \mid l)},$$

which after simplifying can be shown to hold provided only that $s_h^h > s_h^l$. That is, the decision maker is better off under the presumed costless auditing strategy provided that the audit signal $s$ is informative.

Similarly, we must check that the decision maker is better off using both the biased expert's report and her own imperfect signal. This is true provided that:

$$\frac{ps_h^h}{ps_h^h + (1-p)s_h^l} < \frac{ps_h^l}{ps_h^h + (1-p) f(H \mid l) s_h^l} \;,$$

which after simplification becomes the requirement that $f(H|l) < 1$. In other words, provided that the expert's report that regulatory benefits are "high" is at all informative, the decision maker is better off using both her imperfect signal and the decision maker's report than if she regulated based only her report.

To see whether this constraint is indeed satisfied in equilibrium, we derive, finally, the probability $f(H|l)$ that the expert reports "high" when the true state is low which just makes the decision maker indifferent between regulating and not. Under the presumed regulatory strategy of regulating if and only if both the expert and the independent signal report "high," the decision maker is indifferent between regulating and not if and only if:

$$\frac{ps_h^h}{ps_h^h + (1-p)s_h^l f(h \mid l)}\theta_h + \frac{(1-p)s_h^l f(h \mid l)}{ps_h^h + (1-p)s_h^l f(h \mid l)}\theta_l - d = 0. \tag{10}$$

When we rewrite this equality, we have that the optimum probability with which the expert reports *H*, when *l* is the true state, given that the decision maker audits, denoted by $f^{*a}$, is given by:

$$f^{*a} = \frac{ps_h^h(\theta_h - d)}{(1-p)s_h^l(d - \theta_l)} \tag{11}$$

As $f^{*a}$ is bounded from above by 1, it is clear from (11) that as the decision maker's signal becomes perfectly informative, so that $s_h^h = 1$ and $s_h^l = 0$, the expert's report becomes completely uninformative, as the expert always reports "high" no matter what the state (viz., $f(H|l) = 1$). Intuitively enough, if the decision maker's signal becomes completely informative, the expert's report is essentially irrelevant to the regulatory decision, and, in the limit, the decision maker dispenses altogether with an expert report.

Conversely, as the decision maker's signal becomes very uninformative, so that the decision maker is in essence relying solely on the expert, (that is, $s_h^h$ and $s_h^l$ each go to ½), $f(H|l)$ must be sufficiently low to cause the expected, posterior net cost of regulating when the benefits are low to be less than the expected net gain from regulating when the benefits are high. Quite non-intuitively, by relying not only on the expert report but also on her own very poor signal, the decision maker puts the expert in a position where she knows that her report will be determinative, but only if it is very reliable. In other words, what has been shown is that it is precisely when the decision maker has very little accuracy in determining the state of the world that she should make her own determination in addition to that of the expert. This is because of the impact on the incentives of a biased expert who rationally plays the expert reporting game so as to maximize her own payoff.

The other side of the BPE in this case is, as before, the probability of regulation $r^{*a}$ that makes the expert just indifferent between truthful and honest reporting, given that the benefits are observed to be low. This indifference condition requires that:

$$0 = s_h[r(\theta_l - c) + (1-r)(-c)] + (1 - s_h^l)(-c),$$

which simplifies to:

$$r = \frac{c}{s_h^l \theta_l}. \tag{12}$$

When we compare (12) with (5), we see that when the decision maker requires that both her own signal and the expert's report indicate that benefits are high, she must regulate with a higher probability after hearing the expert report high benefits to induce the expert to be willing to provide informative reports (more precisely, to be indifferent between lying and telling the truth in the low benefits state). This is because auditing by the decision maker – her reliance on her own signal in addition to the expert's report – means that there are some occasions when the expert reports high but the decision maker does not regulate. However, by the same token, the equilibrium probability of regulation given the decision maker's signal indicates high benefits is the same as previously.

B. Competing Experts

Another possible response to reports from biased experts is to have both expert types simultaneously report to the decision maker on the state of the world. To identify possible equilibria in such a game, let us first assume that each expert is better off participating than not, given that the other expert participates. Given this assumption, which will be verified as true in equilibrium, each expert must truthfully report in her preferred state – high benefits for the pro-regulation expert, low benefits for the anti-regulation expert. In the unique BPE in mixed strategies, in the "bad" state for each, each expert adopts a probability of truthful reporting that causes the decision maker to be indifferent between regulating and not, given the strategy adopted by the other type of expert. To identify these strategies, first let $g(\Theta|x,y)$ denote the decision maker's posterior probability, given a report $x$ from the pro-regulation expert and $y$ from anti-regulation expert. Let $f^x(i|j)$ and $f^y(i|j)$ for $i = L,H$ and $j = l, h$ denote the report probabilities chosen by, respectively, the pro-regulation $x$ type expert and the anti-regulation $y$ type expert. Now note that there are only three combinations of expert reports $(x,y)$, that can ever be observed in equilibrium: $(h,l)$, $(h,h)$ and $(l,l)$. What cannot be observed is a state where pro-regulation expert reports low benefits and anti-regulation expert reports high, because the pro-regulation expert only reports "low" with positive probability if low is the true state, in which case the anti-regulation expert always reports "low' as well (and vice versa for the case where the true state is high). That expert agreement – the observed reports $(l,l)$ and $(h,h)$ – occurs with positive probability is necessarily true when each expert type adopts a randomized strategy of truthful reporting. It can easily be verified that if the decision maker is told $(l,l)$ then she knows that the state is low and she does not regulate (regulates with probability 0), while if she is told $(h,h)$ then she knows that benefits are high and she always regulates (regulates with probability 1).

The other possible set of reports is $(h,l)$ and this represents the classic case of offsetting expert reports. Given such a pair of reports, the decision maker's posterior probability that the state is "high" is given by:

$$g(h|H,L) = \frac{pf^y(L|h)f^x(H|h)}{pf^y(L|h)f^x(H|h) + (1-p)f^y(L|l)f^x(H|l)} \quad .$$

But the experts truthfully report with probability 1 in their preferred states, and so this posterior probability simplifies somewhat to:

$$g(h|H,L) = \frac{pf^y(L|h)}{pf^y(L|h) + (1-p)f^x(H|l)} . \tag{13}$$

We proceed as before to identify the probability g*(H,L) that makes the decision maker just indifferent between regulating or not, given this report pair. Decision maker indifference requires:

$$g(h \mid H, L)(\theta_h - \theta_l) = k - \theta_l,$$

which, after substituting for $g(h|H,L)$ from (13) and simplifying, becomes:

$$f^x(H \mid l) = \frac{p}{1-p}\left(\frac{\theta_h - k}{k - \theta_l}\right) f^y(L \mid h). \quad (14)$$

Equation solves for the equilibrium probability of false expert reports that make the indifferent between regulating and not. When we compare (14) to (4), we can see that:

$$f^x(H \mid l) < f^*(H \mid l),$$

which means that the pro-regulation expert reports "high" benefits with a lower probability when she can expect to be countered by an anti-regulation expert. The explanation is this: expert competition increases the reliability of an expert's report compared to her report without such competition, because the fact that each expert is sometimes reporting falsely constrains the other expert to report more truthfully, lest the decision maker either prefer to regulate or prefer not to regulate.

Once again, we turn next to solve for the equilibrium probability of regulation given the reports *r(x,y*) that make each expert just indifferent between truthful and versus untruthful reporting, given that the state of the world is the one in which they have an incentive for untruthful reporting. The probability we are after is that which makes both experts just indifferent between truthful and honest reporting. These turn out to be the very same probabilities that we have already indentified. To see this, suppose that the state of regulatory benefits is "low." In this state, the anti-regulation expert always honestly reports *L*. Given this, if the pro-regulation expert also says *L* then regulation occurs with probability 0 and the pro-regulation expert gets a payout of *0*. If the pro-regulation expert says *H* even though benefits are, actually low then her expected payout is $r(H,L)(\Theta_l - c) + (1 - r(H,L)(-c))$. Setting this equal to 0 – her payout from a truthful report, we have that:

$$r^x(H, L) = \frac{c}{\theta_l}. \quad (15)$$

But the right hand side in (15) is precisely equal to the right hand side in (5). The equilibrium probability of regulation necessary for the pro-regulation expert to be indifferent between truth and falsity in her undesired state is the same, regardless of whether we have competing experts. The reason is that in the undesired state of the pro-regulation expert – low benefits – the anti-regulation expert is always truthful. Hence, and somewhat non-intuitively, it is precisely when the pro-regulation expert has an incentive to falsely report – the state of low regulatory benefits – that the presence of the anti-regulation expert has no impact on the the pro-regulation expert's payout from truth versus falsity.

By precisely the same kind of analysis, we find that on the anti-regulator's side, the probability of regulation necessary to make her indifferent between truth and false reports, given high regulatory benefits, is given by:

$$r^{y}(H,L) = \frac{k-c}{k}, \qquad (16)$$

which of course is precisely the same as (7).

Of course, there can be only one probability of regulation given the reports *(H,L)*. If we set (16) equal to (15), we see that such an equilibrium probability of regulation exists iff:

$\frac{k-c}{k} = \frac{c}{\theta_l}$, or, rearranging, if:

$$\frac{\theta_l}{c} = \frac{k}{k-c}. \qquad (17)$$

Roughly speaking, the left hand side in (17) is the pro-regulation expert's gross benefits from the untruthful report divided by her cost of an untruthful report; the right hand side is the anti-regulation expert's gross benefit from an untruthful report divided by her (net) cost of such a report.

There is no particular reason to think that equation (17) will hold. When it does not hold, a BPE equilibrium with competing experts will exist if and only one of the expert is compensated for reporting against her (known) interests. To see this, suppose that the pro-regulation expert is paid an amount *w* whenever she reports that regulatory benefits are low. Solving for and equalizing $r^{y}(H,L)$ and $r^{x}(H,L)$ as above, but with a wage w paid to the truthful pro-regulation expert, we have that:

$$w^{*} = \frac{k(1+c)-c}{k}. \qquad (18)$$

As can be seen from (18), the higher is *k*, the reward to the anti-regulation expert from forestalling regulation, the higher must be the wage paid to a truthful pro-regulation expert who reports against her interests. This is because the higher is *k*, the lower is the equilibrium probability of regulation that makes the anti-regulation expert indifferent between truthful versus false reporting (in the high benefits state). The wage is paid to the pro-regulation expert to ensure that she too is indifferent between truthful and honest reporting at a relatively low probability of regulation.[16]

---

[16] This is analogous to the result on compensation of competing agents reported in Mathias Dewatripont and Jean Tirole, Advocates, 107 *J. Pol. Econ.* 1, 13-14 (1999).

It remains to confirm the earlier assumption that both types of experts will indeed participate in the reporting game. There are again two cases to consider. In the first case, the decision maker will regulate even without an expert report (viz., (1) holds). In this case, the anti-regulation expert would participate, and given participation by the anti-regulation expert, the pro-regulation expert will certainly participate, provided that a wage *w* is paid if necessary to sustain the equilibrium. The inverse holds when the decision maker will not regulate based on prior probabilities alone, with the difference being that the anti-regulation expert might need to be paid a wage for reporting that the state is high.

III. Costly and Potentially Imperfect Expert Information with Experts of Varying Types

The previous parts have analyzed an expert who has perfect information about regulatory benefits but must incur a cost to report to the decision maker. In this part, I consider an expert who must incur a cost to learn about the state of regulatory benefits, and where such learning may be imperfect in the sense that the expert may expend the cost but not actually learn anything more than she knew a priori.[17] Toward this end, assume that the expert's private valuation of regulatory benefits is equal to the true social value $\Theta$ but the expert discounts the true regulatory compliance cost by a factor of $\lambda$ with $0 < \lambda$. The expert perceives a private benefit from regulation of $\lambda\Theta - k$. Suppose also that the regulator's payoff from regulation is simply the actual net benefit of regulation $\Theta - k$. Hence an expert with $\lambda < 1$ may be thought of as being biased in favor of regulation (because she discounts compliance costs), while a regulator with $\lambda > 1$ is biased against regulation (she attaches too much weight to regulatory compliance costs). Compliance costs k are common knowledge; both the regulator and expert know the distribution of $\Theta$, with density $f(\Theta)$ defined on $[0, \bar{\theta}]$, and the expert's type $\lambda$ is known only to the expert, while the regulator knows the distribution $g(\lambda)$ with positive support on $[0, \bar{\lambda}]$.

Consider now the expert's incentives to take costly effort to learn about the true state of regulatory benefits. The expert observes a signal of the true state, with the signal given by $s = \varepsilon\Theta$, where $\varepsilon$ is a random variable with gamma distribution on support $[0, \infty]$ and $E(\varepsilon) = 1$. Suppose that the expert may invest an amount e, at cost $c(e)$ with $c'(e) > 0$ and $c''(e) < 0$, $c(0) = 0$, and that if she invests *e* then with probability $\pi(e)$ the signal s is perfectly informative, and the expert learns the true value of regulatory benefits, while with probability $1 - \pi(e)$ her investment provides no additional information and the signal is uninformative (the expert knows only what she knew before, the density $f(\Theta)$). Upon observing her signal, the expert knows just the probability that it is perfectly accurate versus uninformative. (One may imagine an expert who spends time and effort to survey work on regulatory benefits; the more time and effort she has spent, the greater is her perceived likelihood that her impression of the level of benefits is accurate).

The game then proceeds as follows. The expert first decides whether to invest some amount *e* to generate a probability $\pi(e)$ of learning the true state of regulatory

[17] This is the same simplified learning structure modeled by Dur and Swank (2005) and Dewatripont and Tirole (1999), supra.

benefits, with *e* not observed by the decision maker. After making this decision, the expert then decides whether or not to participate in the expert reporting game, and if she participates, whether to recommend "regulate" or "no regulation." These are the only two reports allowed, and they correspond roughly to the "high" and "low" messages of the game previously considered, where $\Theta$ assumes only two possible values.

Assume first that the regulator will follow the expert's recommendation and regulate if and only if the expert recommends this. Conditions under which this assumption holds will be demonstrated below. After making her effort choice, the expert will get a higher payout from regulation than no regulation and therefore recommend regulation whenever:

$\pi(e)\theta + (1-\pi(e))\mu - \lambda k \geq 0$, or whenever,

$$\theta \geq \frac{\lambda k}{\pi(e)} - \frac{\mu(1-\pi(e)}{\pi(e)} \equiv T(\lambda) \quad . \tag{19}$$

Intuitively enough, inequality (19) says that that the threshold level of signaled regulatory benefits for which the expert recommends regulation increases with the weight that the expert gives to regulatory compliance costs. For sufficiently low such weight, the expert always recommends regulation, and for sufficiently high such weight, the expert never recommends regulation.

Given T(λ), the expert's problem of choosing how much effort to expend in (possibly) becoming informed as to the true state of regulatory benefits is given by:

$$\max_e \left\{ \pi(e)(1-F(T(\lambda)) \left[ \frac{\int_{T(\lambda)}^{\bar{\theta}} \theta f(\theta)d\theta}{1-F(T(\lambda))} - \lambda k \right] + (1-\pi(e))\left[(1-F(T(\lambda))(\mu-\lambda k)\right] - c(e) \right\},$$

which simplifies to:

$$\max_e \left\{ \pi(e)\left[ \int_{T(\lambda)}^{\bar{\theta}} \theta f(\theta)d\theta - (1-F(T(\lambda))\mu) \right] + (1-F(T(\lambda))(\mu-\lambda k) - c(e) \right\}.$$

But this can be rewritten as:

$$\max_e \left\{ \pi(e)\left[ \int_{T(\lambda)}^{\bar{\theta}} (\theta-\mu)f(\theta)d\theta) \right] + (\mu-\lambda k)\int_{T(\lambda)}^{\bar{\theta}} f(\theta)d\theta - c(e) \right\},$$

And with a final rewriting, we have that the expert's problem is given by:

$$\max_e \int_{T(\lambda)}^{\bar{\theta}} \left(\pi(e)(\theta-\mu)+\mu-\lambda k\right) f(\theta)d\theta - c(e). \tag{20}$$

The first order condition for this maximization problem is given by:

$$\pi'(e)\left[\int_{T(\lambda)}^{\bar{\theta}} (\theta-\mu)f(\theta)d\theta\right] - f(T(\lambda))\frac{\partial T}{\partial e}\left[\pi(e)(T(\lambda)-\mu)+(\mu-\lambda k)\right] = c'(e). \tag{21}$$

Consider the left hand side of (21), which gives the expert's marginal benefit of expending effort to possibly learn the true state of regulatory benefits. Because π'(e) > 0, the first term on the left hand side is positive for all values of *T(λ)*.[18] As for the second term on the left hand side, after substitution for *T(λ) = (λk - μ(1- π(e)))/π(e),* we have that the second term is equal to 0. Hence the first order condition reduces to:

$$\pi'(e)\left[\int_{T(\lambda)}^{\bar{\theta}} (\theta-\mu)f(\theta)d\theta\right] = c'(e). \tag{22}$$

If we let e* denote the solution to (22), then by applying the implicit function theorem we derive the comparative statics result that:

$$\frac{\partial e^*}{\partial \lambda} = \frac{-\pi'(e)\left[-T'(\lambda)T(\lambda)f(T(\lambda)) + f(T(\lambda))T'(\lambda)\mu\right]}{\pi''(e)\left[\int_{T(\lambda)}^{\bar{\theta}} \theta f(\theta)d\theta - (1-F(T(\lambda))\mu)\right] - c''(e)}. \tag{23}$$

By the second order condition for e* to constitute a maximum, the denominator in (23) is negative, and so we have that:

$$sign\left(\frac{\partial e^*}{\partial \lambda}\right) = sign\left(\pi'(e)\left[-T'(\lambda)T(\lambda)f(T(\lambda)) + f(T(\lambda))T'(\lambda)\mu\right]\right)$$
$$= sign\left(\pi'(e)T'(\lambda)f(T(\lambda))\left\{\mu - T(\lambda)\right\}\right).$$

But substituting for T(λ) from (19), and using the fact that π'(e) > 0 and T'(λ) = k/π(e) > 0, we have finally that:

---

[18] To see that this is so even for *T(λ) < μ,* suppose to the contrary that for such a *T(λ)* we had $\int_{T(\lambda)}^{\bar{\theta}} (\theta-\mu)f(\theta)d\theta < 0$. Then it would follow that $\int_{0}^{T(\lambda)} (\theta-\mu)f(\theta)d\theta < 0$ and together this inequalities would imply that $\int_0^{\bar{\theta}} (\theta-\mu)f(\theta)d\theta = \mu-\mu < 0$, a contradiction.

$$sign\left(\frac{\partial e^*}{\partial \lambda}\right) = sign\left(\mu - \lambda k\right). \tag{24}$$

From (24), it follows that as e*(λ) reaches its maximum when λ = μ/k: the expert who is indifferent whether to regulate ex ante or not will invest maximum effort to become informed, and effort levels fall off, the more the expert is biased in favor (low λ) or against (high λ) regulation ex ante. This is qualitatively different than in Gur and Swank (2005), where it is the expert whose preferences are perfectly aligned with those of the decision maker who makes the maximum investment to learn the true state of the world. Here, only if the ex ante case for regulation is in perfect balance, in the sense that μ = k, do we have the result that maximum effort to become informed is made by the expert whose preferences are perfectly aligned with the decision maker. When μ > k and the regulator's preference (if she knew μ) would be to regulate ex ante, maximum investment is undertaken by an expert who overweights regulatory compliance costs and in this sense biased against regulation; when μ < k, and the decision maker would prefer not to regulate ex ante, maximum investment is undertaken by an expert who underweights regulatory compliance costs and is in this sense biased in favor of regulation. In this sense, the ideal expert from the point of view of providing information to the decision maker is one whose preferences run against the decision maker's own ex ante preference.

However, relying upon an expert whose preferences run against the decision maker's preferences also involves an obvious cost due to the preference divergence. For the decision maker, there is a tradeoff between expert incentives to become informed about the state of the world and the expert's preferences over outcome. To see this tradeoff most clearly, consider the decision maker's expected payoff from expert investment to become informed. Recall that the decision maker attaches a weight of λ = 1 to regulatory compliance cost *k*. From (20), the decision maker's expected payoff from employing an expert type λ is, for arbitrary λ, given by:

$$V(\lambda) = \pi(e^*(\lambda))\left[\int_{T(\lambda)}^{\bar{\theta}} \theta f(\theta)d\theta - (1 - F(T(\lambda)))\mu\right] + (1 - F(T(\lambda)))(\mu - k), \tag{25}$$

where *T(λ)* is once again as defined by (19).

To see the tradeoffs involved for the decision maker as expert type λ changes, and identify the optimal expert type, λ*, take the derivative of (25) with respect to λ, given by:

$$\begin{aligned}\frac{\partial V}{\partial \lambda} &= \pi(e^*)\left[T'(\lambda)T(\lambda)f(T(\lambda))) + f(T(\lambda))T'(\lambda)\mu\right] - f(T(\lambda))T'(\lambda)(\mu - k) \\ &+\pi'(e^*)\frac{\partial e^*}{\partial \lambda}\left[E(benefit \mid \lambda)\right],\end{aligned} \tag{25}$$

where

$$E(benefit \mid \lambda) \equiv \left[ \int_{T(\lambda)}^{\bar{\theta}} \theta f(\theta) d\theta - (1 - F(T(\lambda))\mu \right].$$

Now substituting for $T(\lambda)$ and rearranging, (25) can be rewritten as:

$$\frac{\partial V}{\partial \lambda} = f(T(\lambda)) \left[ \frac{k(\pi(e^*) - \lambda) + \mu(1 - \pi(e^*))}{\pi(e^*)} \right] + \pi'(e^*) \frac{\partial e^*}{\partial \lambda} \left[ E(benefit \mid \lambda) \right]. \qquad (26)$$

The second term in (26) is the marginal benefit in terms of increased effort to learn the true state of regulatory benefits as a function of $\lambda$; the first term captures the cost of potential preference divergence.

To solve for the optimal expert type, we consider two cases, and use the continuity of (26).

*Case 1:* $\mu < k$. First evaluate (26) at $\lambda = 1$ (an expert whose preferences are identical to those of the decision-maker). At $\lambda = 1$, we have that $\lambda > \mu/k$, and so from (23), we know that $\partial e^*/\partial \lambda < 0$, and so the second term on the RHS of (26) is negative. When $\lambda = 1$, the first term on the RHS of (26) is equal to $f(T(\lambda)) \left[ \frac{(1 - \pi(e^*))(\mu - k)}{\pi(e^*)} \right] < 0$ when $\lambda < k$. Hence when $\mu < k$, we have that $\partial V/\partial \lambda < 0$ at $\lambda = 1$.

Now consider $\lambda = \mu/k < 1$, an expert who makes the maximized effort to become informed as to the state of regulatory benefits, that is, $\partial e^*/\partial \lambda(\lambda = \mu/k) = 0$, and so the second term on the RHS in (26) is also 0. As for the first term on the RHS in (26), when $\lambda = \mu/k$, this term reduces to $f(T(\lambda))(k - \mu) > 0$ for $\mu < k$. Hence, we have that $\partial V/\partial \lambda > 0$ at $\lambda = \mu/k$.

As we have that $\partial V/\partial \lambda > 0$ at $\lambda = \mu/k < 1$ and $\partial V/\partial \lambda < 0$ at $\lambda = 1$ and $V()$ is continuous in $\lambda$, it must be that when $\mu < k$, the optimal expert type $\lambda^*$ is such that $\mu/k < \lambda^* < 1$. The decision maker sacrifices a lower investment by the expert to become informed in order to get an expert whose preference is closer, but not identical, to her own.

*Case 2:* $\mu > k$. As in the previous case, consider first the value of $\partial V/\partial \lambda$ when $\lambda = 1$. As $\mu/k > 1$ in this case, we have from (23) that $\partial e^*/\partial \lambda > 0$ at $\lambda = 1$ and so the second term on the RHS in (26) is positive at $\lambda = 1$. When $\lambda = 1$, we have as already shown that the first term on the RHS of (26) is given by $f(T(\lambda)) \left[ \frac{(1 - \pi(e^*))(\mu - k)}{\pi(e^*)} \right] > 0$ when $\mu > k$.

Considering now $\lambda = \mu/k$, we again know from (23) that the second term on the RHS is 0 because $\partial e^*/\partial \lambda = 0$ when $\lambda = \mu/k$. The first term is (as shown above) given by $f(T(\lambda))(k - \mu)$ when $\lambda = \mu/k$, and when $\mu > k$, we have that $f(T(\lambda))(k - \mu) < 0$. Hence, we have that $\partial V/\partial \lambda < 0$ at $\lambda = \mu/k > 1$.

As we have that $\partial V/\partial\lambda > 0$ at $\lambda = 1 < \mu/k$ and $\partial V/\partial\lambda < 0$ at $\lambda = \mu/k$, we have from the continuity of $V()$ that when $\mu > k$, the optimal expert type $\lambda^*$ is such that $1 < \lambda^* < \mu/k$.

We thus have the following proposition:

*Proposition 1:* The optimal expert type is located between the decision maker's type and the expert type who maximizes investment to learn the true state of regulatory benefits. In this sense, the decision maker accepts expert bias – in the sense of divergence from her own preferences – in order to increase expert incentives to become informed.