JOURNAL OF PUBLIC ECONOMIC THEORY
Department of Economics 415 Calhoun Hall Vanderbilt University Nashville, TN 37235 USA (615) 343-8495 (fax)

Submission Number:JPET-08-00226

Regulatory Policy Design in an Uncertain World

Robert G. Chambers
University of Maryland, College Park

Tigran A. Melkonyan
University of Nevada, Reno

Abstract

The paper examines principal-agent relationships in uncertain environments where beliefs of the contracting parties (the regulator and the firm) are represented by sets of probabilities. In addition to fully characterizing the first-best and the second-best solutions, we examine optimality of zero-risk, fixed-payment schemes and the relationship between the first-best and the second-best solutions. In the second-best world, where the regulator can only contract on the quality of the good, a zero-riskstandard is optimal when the firm has beliefs that are so ambiguous that the firm''s marginal rate of transformation belongs to the set of the firm''s relative probabilities.

The authors thank Peyton Ferrier and Carlos Arnade for helpful comments on an earlier version of this paper. Chambers' research was partially supported by Cooperative Agreement # 58300060058 between the Economic Research Service, USDA and the University of Maryland.
Citation: Robert G. Chambers and Tigran A. Melkonyan, (2012) "Regulatory Policy Design in an Uncertain World", *Journal of Public Economic Theory*, Vol. 12 No. 6 pp. 1081-1107.
Contact: Robert G. Chambers - rchambers@arec.umd.edu, Tigran A. Melkonyan - tmelkonyan@cabnr.unr.edu.
Submitted: October 22, 2008. **Published:** March 07, 2012.

1 Introduction

Government agencies regulating health and environmental hazards "...typically face massive uncertainties about the extent, sometimes even the existence, of those risks..." (Nichols and Zeckhauser, 1986). Empirical evidence and casual introspection suggest, moreover, that regulators and decisionmakers frequently possess information that is consistent with varying theories about the likelihoods of different hazards. Disputes, even among experts, about the likelihood of relevant events are the norm and not the exception. In other words, regulators of health and environmental hazards typically face "scientific uncertainty" in designing their policies. For example, Lee (2005) observes: "Whilst expertise, particularly scientific expertise, is an inescapable part of environmental risk regulation, recent very bruising public encounters with the limitations of science have put an end to the assumption that 'scientific facts' provide adequate authority for regulation" (Lee, 2005, pp. 79-80).

Faced with such uncertainties, regulating agencies frequently resort to regulations based on worst-case scenarios, while rationalizing their regulation strategy on the grounds that it reflects "prudence" or "conservatism" (Nichols and Zeckhauser, 1986; Viscusi and Gayer, 2002). Conservative and prudent regulatory strategies often involve "zero-risk" regulations where an activity is banned or regulatory standards to minimize hazards associated with that activity are mandated.

Perhaps the best known example of a "zero-risk" policy is the Delaney Clause adopted by the US Congress in 1958 as an amendment to the Food, Drugs, and Cosmetic Act of 1938. The Clause calls for a "zero-risk" policy toward food additives classified as carcinogens: "...the Secretary of the Food and Drug Administration shall not approve for use in food any chemical additive found to induce cancer in man, or, after tests, found to induce cancer in animals." (CFR, 2007). It is noteworthy that a number of observers have singled out scientific uncertainty about cancer causation as one of the key factors that originally led to the adaptation of the Delaney Clause (Breyer, 1993, Weisburger, 1996).[1]

[1] In the light of scientific breakthroughs in 1980s and early 1990s it became apparent that the Delaney Clause was based on mistaken scientific notions of cancer causation (Weisburger, 1996). More to the point, the scientific orthodoxy had changed. This has led to a widespread discomfort with the clause (Breyer, 1993). In 1996 a broad consensus on amending the clause was established. The amended version of the Delaney Clause gave governmental agencies flexibility to implement the law at essentially the prevailing acceptable-risk standard for non-carcinogens, which the legislative language calls "reasonable certainty of no

Another example of a "zero-risk" policy is the U.S. Safe Drinking Water Act, which mandates setting the maximum contaminant levels goals (MCLGs) and the maximum contaminant levels (MCLs) of the nation's water supply. The act states that the MCLGs "shall be set at the level at which no known or anticipated adverse effects on the health of persons occur and which allows an adequate margin of safety."(42 U.S.C. Sec.300g-l(b)(3)).

The European Union has also frequently employed such conservative regulatory policies. Most noteworthy are its policies grounded in the precautionary principle (PP). The PP has been used by the EU to mandate zero-risk policies for a number of health and environmental hazards. For example, between October, 1998 and May 2004, the EU did not authorize any genetically modified organisms for release in the European Union (CEC 2000).

This paper examines optimal regulatory design (more generally principal-agent relationships) in uncertain (*ambiguous)* environments that arise because of the presence of scientific uncertainty about potential health or environmental hazards. A particular focus of the paper is to determine whether there exist conditions under which the presence of scientific uncertainty merits the imposition of a "zero-risk" regulatory stance as part of an optimally designed policy.

In considering these issues, our paper necessarily departs significantly from much of the literature on optimal regulation of health and environmental hazards. With very few exceptions, existing analyses of optimal regulatory policy design (more generally, principal-agent relationships) assume that involved parties have expected-utility preferences and that there is no scientific uncertainty about potential hazards. Thus, beliefs about the state of the world are represented by unique additive (objective or subjective) probabilities, common to all parties, over the relevant outcomes. In this setting, a "zero-risk" regulation is an extreme outcome that will be optimal only under the most stringent assumptions.

The applicability of these models to situations involving scientific uncertainty, however, is questionable. Not only must economic agents be able to form subjective probabilities, these subjective probabilities must be shared by all actors in a regulatory setting. It is hard to imagine such agreement occurring when there exists significant scientific uncertainty about

harm." Pesticide use was removed from the Delaney Clause in 1996 by an amendment to Title IV of the Food Quality Protection Act.

the presence, nature, or probability of a particular environmental or health hazard.

Ellsberg (1961) argued that, when faced with an uncertain decision environment, individuals exhibit behavior sensitive to the weight of evidence about probabilities (the famous "urn" examples). Such behavior directly contradicts both objective and subjective expected utility theory, and, if descriptive of reality, renders expected-utility theory (more generally probabilistically sophisticated behavior) inappropriate for evaluating situations involving uncertainty. Ellsberg-type behavior has been repeatedly validated in the experimental and empirical literatures.[2]

Therefore, to accommodate the potential presence of scientific uncertainty in a regulatory setting, we model individual attitudes towards uncertainty as represented by the Gilboa-Schmeidler (1989) maximin expected-utility (MMEU) model. Beliefs of a decision-maker for this preference structure are represented by a *set* of probability distributions in contrast to a unique probability distribution for a probabilistically sophisticated decision-maker (Machina and Schmeidler, 1992).

For concreteness sake, our model is cast in a particular regulatory environment. An incentive scheme established by the regulator (the principal) affects the choice of the stochastic quality of the good produced by the firm (the agent). This setup reflects regulation of product safety such as foods, chemicals, and other potentially hazardous products.

In what follows, we first specify the model, and then we treat the first-best case. After developing the first-best results, we study the optimal incentive scheme established by an informationally constrained regulator. A characterization of optimal informationally-constrained contracts between the principal and the agent follows. These results are then applied to determine necessary and sufficient conditions for the optimality of "zero-risk" regulation. We then demonstrate that our framework is particularly suitable for examining food safety regulations. In the following section we relate our work to other studies that have considered the optimal design of principal-agent relationships in an uncertain setting, and the last section concludes.

[2] Camerer (1995) provides an accessible introduction to the empirical literature on the Ellsberg Paradox.

2 The Model

There are two parties in the model, a regulator and a firm. The firm produces a single unit of output of uncertain quality, where the uncertain quality reflects either the potential presence of a foodborne health hazard or an environmental contaminant. Uncertainty is represented by the state space $\Omega = \{1, 2\}$.[3] Random variables thus correspond to elements of $\mathbb{R}^2$. Hence, the *ex ante* uncertain quality of the output is denoted by $\mathbf{z} = (z_1, z_2)$. The firm's *ex ante* minimal cost of producing output with uncertain quality, $\mathbf{z}$, is

$$c(\mathbf{z}) = \alpha z_1 + z_2, \ \alpha > 0.$$

The only domain restriction on z_1 and z_2 is that they belong to $\mathbb{R}_+$. Otherwise, they are subject to choice by the firm producing the uncertain quality. This cost-structure, which is state-allocable in the sense of Chambers and Quiggin (2000), assumes that uncertain quality is perfectly substitutable across states of Nature from the producer's perspective. Thus, it maintains that producers can transform quality across the different states of Nature at a constant cost. The assumption of a constant marginal cost for each state-contingent quality simplifies the mathematical presentation, but can be easily generalized, following Quiggin and Chambers (1998), to any constant returns to scale production structure.

The regulator cannot observe the realization $s \in \Omega$, but it can costlessly observe the *ex post* (realized) quality level. In terms of the uncertain quality $\mathbf{z} = (z_1, z_2)$, this assumption means that only the firm knows $\mathbf{z} = (z_1, z_2)$ *a priori.* However, both the firm and the regulator can observe costlessly the *ex post* realization of $\mathbf{z}$. We assume that this quality level is contractible, and that the regulator can enforceably condition a monetary transfer (or penalty) to the firm on this quality. Following Quiggin and Chambers (1998) and Chambers and Quiggin (2005), we represent the payment scheme from the regulator to the firm as follows: when the firm realizes an *ex post* quality of, say, z, it receives a payment of y_1 if $z = z_1$, a payment of y_2 if $z = z_2$, and an arbitrarily large negative payment otherwise. This mechanism corresponds to actual regulations that require testing of the final product and

[3] Our arguments can be extended from the two-state case to the S-state case and to an atomless continuum by using the methods developed in Chambers and Quiggin (2005) and Racionero and Quiggin (2006), respectively.

impose penalties that vary with the observed level of quality (for example, the pathogen count of a tested product for the case food safety regulations).

The timing of the interaction between the regulator and the firm is as follows: The regulator offers a contract structure to the firm that specifies that the firm will receive y_s if z_s ($s \in \Omega$) is realized and an arbitrarily large negative payment otherwise; the firm then makes its *ex ante* production choice; Nature resolves uncertainty by making a choice from Ω; the *ex post* (realized) quality is observed, and the regulator makes the payment to the firm.

We are particularly interested in "zero-risk" standards based on worst-case scenarios. In our model, such a standard is most naturally interpreted as a fixed-payment structure that induces choice of minimum quality invariant to the realized state of nature (Chambers and Quiggin, 2005). Because the quality is invariant to the realized state of nature, there is literally no uncertainty about the outcome and, thus, "zero-risk" for an arbitrary prior probability.

It may seem that this notion of a "zero-risk" standard does not correspond directly to the types of prohibitions on sales that are involved in such "zero-risk" regulations as the Delaney Clause. However, in our treatment, the quantity of the good marketed is fixed, and the presence of a hazard is measured (monotonically) by the magnitudes of z_1 and z_2. Thus, if a certain quality level is deemed to be unacceptable as a health outcome, the natural regulatory response is to design the payment scheme in such a fashion that the *ex ante* choices of the firm preclude that quality level from emerging. Hence, by setting the *ex ante* quality at a relatively high and nonstochastic level, the contract between the regulator and the firm effectively embargoes the sales of potentially hazardous items.

Both the regulator and the firm have maximin expected utility preferences (Gilboa and Schmeidler, 1989). The regulator's preferences over quality of the output and the payments are represented by the functional:

$$W^P(\mathbf{z}, \mathbf{y}) = \min\left\{\pi_1\left[u(z_1) - y_1\right] + \pi_2\left[u(z_2) - y_2\right] : \boldsymbol{\pi} \in P\right\},$$

where $u(\cdot)$ is a strictly increasing, smooth, strictly concave function, and the set $P =$

$\left\{(\pi_1, 1-\pi_1) : \pi_1 \in \left[\underline{\pi}_1^P, \bar{\pi}_1^P\right] \subset (0,1)\right\}$ represents the regulator's beliefs about the true underlying state of Nature. The limiting cases where $\underline{\pi}_1^P = 0$ and/or $\bar{\pi}_1^P = 1$ are straightforward and are, therefore, omitted. For any given prior the regulator is risk-averse in the sense that it always prefers a nonstochastic quality to a stochastic quality with mean (for that prior) equalling the nonstochastic quality. In this sense, the regulator is presumed to reflect Society's aversion to the health hazard associated with the variable quality.

The firm's preferences are represented by the functional

$$W^A(\mathbf{z}, \mathbf{y}) = \min\{\pi_1 y_1 + \pi_2 y_2 : \boldsymbol{\pi} \in A\} - c(\mathbf{z}),$$

where $A = \left\{(\pi_1, 1-\pi_1) : \pi_1 \in \left[\underline{\pi}_1^A, \bar{\pi}_1^A\right] \subset (0,1)\right\}$ represents the firm's beliefs about the state of the world. Thus, the firm is assumed to be risk-neutral with regard to both the quality and the associated payment over the two states. Abstracting from risk considerations in examining the firm's optimal reaction to contracts proposed by the regulator enables us to focus more narrowly on the role that uncertainty, and beliefs about that uncertainty, play in determining optimal contract form.

The probabilities in P and A can be thought of as representing at least two factors: the relevant decision-maker's information on the possible probability distributions and his or her *degree of confidence* in the existing theories surrounding these probability distributions. This interpretation of beliefs can be traced back to Ellsberg (1961). So, for example, if there are several competing hypotheses about the stochastic structure that characterizes the food-borne or environmental health hazard, but the principal is convinced that only one is truly valid, then P would be a singleton. In that case, the principal's preference structure is standard subjective expected utility. Hence, expected utility is a special case of our model. Conversely, if the principal had no confidence in any of the theories the set P could be quite large, and, in the extreme, could encompass the entire unit interval. The length of the interval which the decisionmaker will entertain as possible probabilities (for example, for the principal $\bar{\pi}_1^P - \underline{\pi}_1^P$) will be referred to as the decision-maker's *degree of imprecision* (Walley, 1991) in what follows.

It is common practice in the analysis of regulatory design (and principal-agent models) to assume that all parties share common beliefs about the uncertain nature of the world. In a situation, such as the one modeled here, where beliefs are inherently subjective, there is no *a priori* reason to expect that these beliefs should always coincide. This seems particularly true when enough scientific uncertainty exists to inhibit the formulation of singleton subjective probabilities. Thus, while our model allows A and P to be common, it does not require it.

We do assume, however, that both the stochastic technology and the belief structures of the firm and the regulator are common knowledge. Generalizing the model to allow for the case where either or both of these are hidden knowledge greatly complicates the regulatory-design (principal-agent) problem, and thus necessarily distracts attention from the main focus of this paper, which is the impact that the presence of scientific uncertainty has on regulatory design.

3 The First Best

As a point of reference, we first consider the regulator's first-best problem. Formally, it is represented as

$$\max_{\mathbf{y},\mathbf{z}} \left\{ W^P(\mathbf{z},\mathbf{y}) : W^A(\mathbf{z},\mathbf{y}) \geq \bar{u} \right\},$$

where $\bar{u}$ denotes the firm's reservation utility. By the structure of the regulator's and the firm's preferences, it follows easily that the first best is characterized by $W^A(\mathbf{z},\mathbf{y}) = \bar{u}$. The Lagrangean expression associated with this problem is, therefore,

$$L(\mathbf{z},\mathbf{y},\mu) = W^P(\mathbf{z},\mathbf{y}) - \mu\left(\bar{u} - W^A(\mathbf{z},\mathbf{y})\right), \tag{1}$$

where μ denotes the Lagrangean multiplier. Given the structure of the preferences and the technology, the Lagrangean is concave in $(\mathbf{z},\mathbf{y})$. When a maximum exists, denote

$$(\mathbf{z}^*,\mathbf{y}^*,\mu^*) \in \arg\max_{\mathbf{z},\mathbf{y},\mu} L(\mathbf{z},\mathbf{y},\mu),$$

$$P^* = \arg\min\left\{\pi_1\left[u(z_1^*) - y_1^*\right] + \pi_2\left[u(z_2^*) - y_2^*\right] : \boldsymbol{\pi} \in P\right\},$$

and

$$A^* = \arg\min \{\pi_1 y_1^* + \pi_2 y_2^* : \boldsymbol{\pi} \in A\}.$$

The conditions that must be satisfied for an *interior* production and payment equilibrium ($z_1^*, z_2^*, y_1^*, y_2^* > 0$), therefore, include

$$\begin{aligned} u'(z_1^*) \min \{\pi_1 : \boldsymbol{\pi} \in P^*\} - \mu^* \alpha &= 0, \\ u'(z_2^*) \min \{\pi_2 : \boldsymbol{\pi} \in P^*\} - \mu^* &= 0, \\ -\max \{\pi_1 : \boldsymbol{\pi} \in P^*\} + \mu^* \min \{\pi_1 : \boldsymbol{\pi} \in A^*\} &= 0, \\ -\max \{\pi_2 : \boldsymbol{\pi} \in P^*\} + \mu^* \min \{\pi_2 : \boldsymbol{\pi} \in A^*\} &= 0, \end{aligned} \tag{2}$$

and

$$\bar{u} - (\min \{\pi_1 y_1^* + \pi_2 y_2^* : \boldsymbol{\pi} \in A\} - \alpha z_1^* - z_2^*) = 0.$$

Conditions (2) have a familiar and intuitive interpretation. α represents the firm's (constant) marginal rate of transformation between quality produced in state 1 and state 2. The expression

$$\frac{u'(z_1^*) \min \{\pi_1 : \boldsymbol{\pi} \in P^*\}}{u'(z_2^*) \min \{\pi_2 : \boldsymbol{\pi} \in P^*\}},$$

represents the regulator's equilibrium marginal rate of substitution between qualities of the good in the two states. Hence, the first two expressions in (2) require that the regulator's marginal rate of substitution is equated to the firm's marginal rate of transformation,

$$\frac{u'(z_1^*) \min \{\pi_1 : \boldsymbol{\pi} \in P^*\}}{u'(z_2^*) \min \{\pi_2 : \boldsymbol{\pi} \in P^*\}} = \alpha.$$

This is the condition (production efficiency) that would emerge if the regulator could forego contracting with the firm over quality and directly undertake planning and implementing the production process himself or herself. In the first best, where information asymmetries play no role, no losses arise from allowing the firm to produce instead of having the regulator control production directly.

μ^* is the marginal benefit to the regulator of a one-unit decrease in the firm's reservation utility, and thus $\mu^* \min \{\pi_s : \boldsymbol{\pi} \in A^*\}$ represents the marginal benefit to the regulator of a small marginal change in payment y_s to the firm while $\max\{\pi_s : \boldsymbol{\pi} \in P^*\}$ represents the (direct) marginal welfare cost to the regulator of such a payment. For there to be an interior payments equilibrium, these terms, as usual (exchange efficiency), must equal one another.

From these conditions we can characterize the conditions under which the regulator requires the firm to produce quality nonstochastically in the first best. Because a "zero-risk" regulatory stance requires nonstochastic production of quality, this also characterizes the case where a "zero-risk" policy could potentially be optimal. We have: (Proofs not presented directly in the text are contained in the Appendix.)

Proposition 1 *The first-best is characterized by* $z_1^* = z_2^* > 0$ *and* $y_1^* = y_2^*$ *if and only if one of the following conditions is satisfied:* $\frac{\underline{\pi}_1^P}{\left(1-\bar{\pi}_1^P\right)} = \alpha$; $\frac{\underline{\pi}_1^P}{\left(1-\underline{\pi}_1^P\right)} = \alpha$; $\frac{\bar{\pi}_1^P}{\left(1-\bar{\pi}_1^P\right)} = \alpha$; *or* $\frac{\bar{\pi}_1^P}{\left(1-\underline{\pi}_1^P\right)} = \alpha$.

These conditions require that one of the regulator's potential marginal rates of substitution between money income in state 1 and state 2 be equated to the producer's marginal rate of transformation between quality in state 1 and state 2. Hence, first-best equilibrium can involve a zero-risk policy if and only if the firm's cost structure is *not inherently risky* in the sense of Chambers and Quiggin (2000) for one of these potential marginal rates of substitution for the regulator. A technology is defined to be *inherently risky at z for the prior π* if it is more costly to produce the mean of $\mathbf{z}$ (for the prior π) with certainty than it is to produce $\mathbf{z}$.

If the technology is not inherently risky for the relevant probabilities, it is (at least weakly) cheaper to produce the mean quality (calculated for that prior) with certainty than it is to produce an output with stochastic quality. Because its preferences are concave over quality outcomes, for any given prior, the principal always prefers the mean quality (calculated for that prior) to any uncertain quality. For any given prior, the principal is, therefore, risk-averse in its usual sense. Therefore, other things equal, in the first-best it would forego that certain quality only if there were some cost advantage associated with producing the uncertain quality which the principal, in turn, could capture and use to compensate itself for its greater risk exposure (again for that prior). However, if the technology were not inherently risky, then, at the margin, there is no cost reduction associated with the firm producing a stochastic quality. Hence, the regulator designs a nonstochastic "zero-risk" payment-quality scheme.

In interpreting, Proposition 1, one should notice that the set of cost parameters, α, for which the proposition is satisfied will have measure zero in the parameter space. Hence, in that sense, it is an extreme requirement, and for most model specifications, one would not expect the first-best policy to involve a "zero-risk" payment-quality scheme.

From (2), there is a strictly positive payment to the firm in state s only if

$$\mu^* = \frac{\max\{\pi_s : \boldsymbol{\pi} \in P^*\}}{\min\{\pi_s : \boldsymbol{\pi} \in A^*\}},$$

from which we can conclude that there can be a strictly positive payment in both states only if

$$\frac{\max\{\pi_1 : \boldsymbol{\pi} \in P^*\}}{\max\{\pi_2 : \boldsymbol{\pi} \in P^*\}} = \frac{\min\{\pi_1 : \boldsymbol{\pi} \in A^*\}}{\min\{\pi_2 : \boldsymbol{\pi} \in A^*\}}.$$

The expression on the left of this equality is the regulator's marginal rate of substitution between income paid out in the two states. The expression on the right is the firm's marginal rate of substitution between income received in the two states. This is a standard exchange efficiency condition that must be satisfied by any interior Paretian equilibrium. Notice, however, that it tells us something important about the relationship between the regulator's beliefs about the state of the world and those of the firm. Rewriting this expression as

$$\frac{\max\{\pi_1 : \boldsymbol{\pi} \in P^*\}}{1 - \min\{\pi_1 : \boldsymbol{\pi} \in P^*\}} = \frac{\min\{\pi_1 : \boldsymbol{\pi} \in A^*\}}{1 - \max\{\pi_1 : \boldsymbol{\pi} \in A^*\}}$$

shows that there can be a strictly interior payment scheme only if the regulator and the firm's belief structures overlap. In fact, we have:

Proposition 2 *The first-best is characterized by strictly positive payments in both states only if* $P \cap A \neq \varnothing$.

Thus, while the regulator and the firm can have divergent beliefs about the true state of the world, these beliefs cannot be too divergent if the first best involves a positive reward in both states of the world.[4] In the words of Billot et al. (2000), the regulator and the

[4] Proposition 2 is similar in spirit to the characterizations of Pareto optimal allocations in exchange

firm must be between agreeing and disagreeing about their beliefs. While their beliefs do not have to be the same, they must overlap, and the regulator and the firm must share at least one prior in common. And, unless the regulator and the firm can agree on at least one prior, the regulator will always award the firm a zero payment in one state of the world (or a negative payment if that is consistent with individual rationality). It follows from the optimality conditions for the first-best problem that in the case where $P \cap A = \varnothing$, a zero payment will be made to the firm in the state that solves:

$$\min \left\{ \frac{\max \{\pi_1 : \boldsymbol{\pi} \in P^*\}}{\min \{\pi_1 : \boldsymbol{\pi} \in A^*\}}, \frac{\max \{\pi_2 : \boldsymbol{\pi} \in P^*\}}{\min \{\pi_2 : \boldsymbol{\pi} \in A^*\}} \right\}.$$

4 The Second Best

The regulator's second-best optimization problem is

$$\max_{\mathbf{y},\mathbf{z}} \left\{ W^P(\mathbf{z},\mathbf{y}) : \mathbf{z} \in \arg\max W^A(\mathbf{z},\mathbf{y}), W^A(\mathbf{z},\mathbf{y}) \geq \bar{u} \right\},$$

where the first constraint is the *incentive compatibility* (IC) constraint while the second constraint is the *individual rationality* (IR) *constraint.*

The firm's and the regulator's preference functionals can be written, respectively, as

$$W^A(\mathbf{z},\mathbf{y}) = \begin{cases} \underline{\pi}_1^A y_1 + \left(1 - \underline{\pi}_1^A\right) y_2 - \alpha z_1 - z_2, & \text{when } y_1 \geq y_2 \\ \bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2, & \text{when } y_1 < y_2 \end{cases} \tag{3}$$

and

$$W^P(\mathbf{z},\mathbf{y}) = \begin{cases} \underline{\pi}_1^P [u(z_1) - y_1] + (1 - \underline{\pi}_1^P) [u(z_2) - y_2], & \text{when } u(z_1) - y_1 \geq u(z_2) - y_2 \\ \bar{\pi}_1^P [u(z_1) - y_1] + (1 - \bar{\pi}_1^P) [u(z_2) - y_2], & \text{when } u(z_1) - y_1 < u(z_2) - y_2 \end{cases}. \tag{4}$$

Given the structure of contracts detailed above, $\mathbf{z} \in \arg\max W^A(\mathbf{z},\mathbf{y})$ requires that the following incentive compatibility constraints be satisfied if $y_2 > y_1$ (Quiggin and Chambers,

economies with ambiguity averse decisionmakers. In these models, a nonempty intersection of the decisionmakers' sets of probabilities is often necessary and sufficient for full insurance (see, e.g., Billot et al., 2000).

1998):

$$
\begin{aligned}
(IC_1) \quad & \bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 && \geq && y_1 - \alpha z_1 - z_1; \\
(IC_2) \quad & \bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 && \geq && y_2 - \alpha z_2 - z_2; \\
(IC_3) \quad & \bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 && \geq && \bar{\pi}_1^A y_2 + \left(1 - \bar{\pi}_1^A\right) y_1 - \alpha z_2 - z_1.
\end{aligned}
\tag{5}
$$

If $y_2 \leq y_1$, the incentive compatibility constraints are:

$$
\begin{aligned}
(IC_1)' \quad & \underline{\pi}_1^A y_1 + \left(1 - \underline{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 && \geq && y_1 - \alpha z_1 - z_1; \\
(IC_2)' \quad & \underline{\pi}_1^A y_1 + \left(1 - \underline{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 && \geq && y_2 - \alpha z_2 - z_2; \\
(IC_3)' \quad & \underline{\pi}_1^A y_1 + \left(1 - \underline{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 && \geq && \underline{\pi}_1^A y_2 + \left(1 - \underline{\pi}_1^A\right) y_1 - \alpha z_2 - z_1.
\end{aligned}
\tag{6}
$$

Following Quiggin and Chambers (1998), it is straightforward to verify that incentive compatibility requires that $(y_1 - y_2)(z_1 - z_2) \geq 0$. This can be established as follows. Consider, for example, the case when $y_2 > y_1$. (IC_1) implies

$$\left(1 - \bar{\pi}_1^A\right)(y_2 - y_1) \geq z_2 - z_1,$$

whence $z_2 - z_1 > 0 \Rightarrow y_2 - y_1 > 0$. Similarly, (IC_2) implies:

$$\bar{\pi}_1^A (y_1 - y_2) \geq \alpha (z_1 - z_2)$$

whence $z_1 - z_2 > 0 \Rightarrow y_1 - y_2 > 0$. From the observation that incentive compatibility requires that $(y_1 - y_2)(z_1 - z_2) \geq 0$, we, therefore, observe that if $z_2 - z_1 > 0$, then (5) are relevant, while when $z_1 - z_2 > 0$ (6) are relevant.

It follows directly from (5) and (6) that

Proposition 3 *a) If* $\alpha \geq \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$*, then only quality vectors* (z_1, z_2) *with* $z_2 \geq z_1$ *are implementable; b) If* $\alpha \leq \frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}$*, then only quality vectors* (z_1, z_2) *with* $z_2 \leq z_1$ *are implementable; c) If* $\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A} \leq \alpha \leq \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$*, then only quality vectors* (z_1, z_2) *with* $z_2 = z_1$ *are implementable.*

Parts a) and b) of Proposition 3 show that stochastic quality vectors are implementable only if they are inherently risky for the prior that the firm will use in assessing stochastic outcomes. Part c) of the proposition shows that if the firm's belief structure is such that the technology is not inherently risky for any of the priors that it will use in evaluating stochastic outcomes, then only contracts involving zero quality risk are implementable.

Visually, the condition, $\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A} \leq \alpha \leq \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$, associated with part c) requires that the firm's isocost curve supports the firm's objective function in the neighborhood of the sure thing. Hence, if the firm were the residual claimant for the output produced, it would rationally choose zero quality risk under these circumstances. Because a nonstochastic quality can always be made both incentive compatible and individually rational by an appropriate choice of a nonstochastic payment, in this case the characteristics of the technology when combined with the firm's belief structure eliminate the incentive problems that a regulator would usually face. Both the firm and the regulator, at least weakly, prefer a nonstochastic quality, the regulator because he or she is averse to quality risk for any given prior, the firm because a nonstochastic quality has cost advantages.

In what follows we only study the case where $\alpha > \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$, and hence $z_2 \geq z_1$, in detail. The analysis for the case $\alpha < \frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}$ is similar, and, therefore, is left to the reader. The actual design of an optimal policy when $\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A} \leq \alpha \leq \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$, although of particular interest to our study, is trivial. The regulator demands the constrained socially optimal "zero-risk" quality from the regulated firm, and, in return, it grants the firm a nonstochastic payment that ensures that the firm is just left indifferent to its reservation utility.

Given the above observations, the regulator's optimization problem when $\alpha > \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$ can be written as

$$\max_{\mathbf{y},\mathbf{z}} \left\{ \begin{array}{c} \min\left\{\pi_1\left[u(z_1)-y_1\right]+(1-\pi_1)\left[u(z_2)-y_2\right] : \pi_1 \in \left[\underline{\pi}_1^P, \bar{\pi}_1^P\right]\right\}; \\ \bar{\pi}_1^A y_1 + \left(1-\bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 \geq y_1 - \alpha z_1 - z_1; \\ \bar{\pi}_1^A y_1 + \left(1-\bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 \geq y_2 - \alpha z_2 - z_2; \\ \bar{\pi}_1^A y_1 + \left(1-\bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 \geq \bar{u} \end{array} \right\}. \tag{7}$$

In writing (7), we have used the fact (shown formally in the proof of Proposition 3) that under the stated conditions, IC_3 is redundant. Figure 1, therefore, represents the set of incentive schemes that are individually rational and incentive compatible for given $\mathbf{z}$ with $z_2 > z_1$.

Payment schemes, (y_1, y_2), to the northeast of line IR satisfy the individual rationality constraint

$$\bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 \geq \bar{u}.$$

Payment schemes to the northwest of line IC_1 satisfy incentive compatibility constraint (IC_1), while incentive schemes to the southeast of line IC_2 satisfy incentive compatibility constraint (IC_2). The relative positioning of IC_1 and IC_2 in Figure 1 is established from the fact that together they imply

$$\left(1 - \bar{\pi}_1^A\right)(y_2 - y_1) \geq z_2 - z_1 \geq \frac{\bar{\pi}_1^A}{\alpha}(y_2 - y_1),$$

and the assumption that $z_2 > z_1$. When $z_2 = z_1$, IC_1 and IC_2 collapse to a single constraint which is represented by the 45^0 line emanating from the origin. The shaded area in Figure 1 thus represents the payment schemes that satisfy the constraints in problem (7), and are thus implementable for fixed (z_1, z_2).

Because the regulator's utility is strictly decreasing in $\mathbf{y} = (y_1, y_2)$, inspection of Figure 1 shows that for any $z_2 \geq z_1$ the regulator will choose an implementable payment scheme on segment AB (for $z_1 = z_2$ the regulator will choose the incentive scheme at the intersection of curve IR and the 45^0 line). Thus, the individual rationality constraint is binding at the optimum.

Lemma 4 *The optimum is characterized by a binding individual rationality constraint:*

i) $\bar{\pi}_1^A y_1^o + \left(1 - \bar{\pi}_1^A\right) y_2^o - \alpha z_1^o - z_2^o = \bar{u}$ *when* $\alpha > \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$

ii) $\underline{\pi}_1^A y_1^o + \left(1 - \underline{\pi}_1^A\right) y_2^o - \alpha z_1^o - z_2^o = \bar{u}$ *when* $\alpha < \frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}$.

iii) $\pi_1^A y_1^o + \left(1 - \pi_1^A\right) y_2^o - \alpha z_1^o - z_2^o = \bar{u}$ *for some* $\pi_1^A \in \left[\underline{\pi}_1^A, \bar{\pi}_1^A\right]$ *when* $\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A} \leq \alpha \leq \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$

By imposing the regulator's indifference curve over payments schemes (not drawn) on the set of implementable contracts in Figure 1, one observes that potentially any payment scheme on line segment AB in Figure 1 can prevail in equilibrium. Thus, there are three qualitatively different cases that can emerge in equilibrium. IC_1 and the individual rationality constraint are binding in the optimum. IC_2 and the individual rationality constraint are binding in the optimum. And finally, only the individual rationality constraint is binding in the optimum. We consider each of these cases in turn by hypothesizing that a specific case materializes in the equilibrium and then characterizing the optimal contract for that case. Because the style of argument is somewhat repetitive, in what follows we focus our exposition on the first case (IC_1 and the individual rationality constraint are binding) while relegating a more detailed treatment of the remaining two cases to the Appendix.

Suppose that IC_1 and the individual rationality constraint are binding in the optimum. Then, the payment schemes that are implementable for (z_1, z_2) are given by:

$$\begin{aligned} \bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 &= y_1 - \alpha z_1 - z_1, \\ \bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 &= \bar{u}. \end{aligned}$$

Solving gives:

$$\begin{aligned} y_1 &= \bar{u} + (1+\alpha) z_1, \\ y_2 &= \bar{u} + (1+\alpha) z_1 + \frac{\delta}{(1-\bar{\pi}_1^A)}. \end{aligned}$$

where $\delta \equiv z_2 - z_1$. By substituting these expressions into the regulator's objective function we obtain a concentrated objective function that can be used to characterize the equilibrium contract structure by the optimal choice of z_1 and δ with $z_2 = z_1 + \delta$, $\delta \geq 0$

$$V(z_1, \delta) \equiv \min \left\{ \pi_1 [u(z_1)] + (1-\pi_1) \left[u(z_1+\delta) - \frac{\delta}{(1-\bar{\pi}_1^A)} \right] : \pi_1 \in \left[\underline{\pi}_1^P, \bar{\pi}_1^P \right] \right\} - (\bar{u} + (1+\alpha) z_1).$$

This concentrated objective function is concave in (z_1, δ). Hence, any local solution must also be a global solution to this problem. Denote:

$$(z_1^o, \delta^o) \in \underset{(z_1,\delta)}{\arg\max} V(z_1, \delta),$$

and

$$P^o \equiv \arg\min \left\{ \pi_1 u(z_1^o) + (1-\pi_1) \left[u(z_1^o + \delta^o) - \frac{\delta^o}{(1-\bar{\pi}_1^A)} \right] : \pi_1 \in \left[\underline{\pi}_1^P, \bar{\pi}_1^P \right] \right\}.$$

Optimality requires that the (one-sided) directional derivative of $V(z_1, \delta)$ (Rockafellar, 1970; Clarke, 1983), whose existence is guaranteed by concavity and denoted by $V'\left(z_1, \delta; \hat{z}_1, \hat{\delta}\right)$, be non-positive for all possible directions $\left(\hat{z}_1, \hat{\delta}\right)$. Calculation shows that

$$V'\left(z_1^o, \delta^o; \hat{z}_1, \hat{\delta}\right) = \min\left\{\pi_1 \left[u'(z_1^o)\right] \hat{z}_1 + (1-\pi_1)\left[u'(z_1^o+\delta^o)\left(\hat{z}_1+\hat{\delta}\right) - \frac{\hat{\delta}}{(1-\bar{\pi}_1^A)}\right] : \pi_1 \in P^o\right\} - (1+\alpha)\,\hat{z}_1 \tag{8}$$

From (8) we obtain

Proposition 5 *Suppose that* $\alpha > \frac{\bar{\pi}_1^A}{(1-\bar{\pi}_1^A)}$ *and* (IC_1) *and* (IR) *are binding at the optimum, then at the optimum* $z_2^o > z_1^o$ *and* $y_2^o > y_1^o$*. A "zero-risk" payment-quality scheme is not optimal.*

Thus, the optimal contract structure is characterized by a stochastic payment scheme $(y_2^o > y_1^o)$ and a stochastic quality vector $(z_2^o > z_1^o)$. This is a direct consequence of condition $\alpha > \frac{\bar{\pi}_1^A}{(1-\bar{\pi}_1^A)}$, which implies that the technology is inherently risky for this prior and $z_2 \geq z_1$. It is explained as follows. Suppose to the contrary that the regulator imposes a "zero-risk" payment-quality scheme on the firm. Because the technology is inherently risky for this prior, producing the associated nonstochastic quality is (at least weakly) more costly than producing a stochastic quality with the same mean (for this prior). Hence, cost savings are available by moving from a nonstochastic to a stochastic quality. And, because all incentive compatibility constraints are satisfied at a nonstochastic payment-quality contract, the regulator can marginally increase her payoff, without (locally) violating any of the individual rationality and incentive compatibility constraints, by moving away from a "zero-risk" payment-quality scheme towards a stochastic payment-quality scheme.

Taking $\left(\hat{z}_1, \hat{\delta}\right) = (1, 0)$ and $\left(\hat{z}_1, \hat{\delta}\right) = (0, 1)$ in (8) yields, respectively, the following first order conditions for the optimization of the regulator's concentrated objective function:

$$\begin{aligned} \min\left\{\pi_1 u'(z_1^o) + (1-\pi_1) u'(z_2^o) : \pi_1 \in P^o\right\} &= (1+\alpha), \\ u'(z_2^o) &= \frac{1}{(1-\bar{\pi}_1^A)}, \end{aligned} \tag{9}$$

where we have used the fact established in Proposition 5 that $\delta^o > 0$. The left-hand side of the first expression in (9) represents the change in the regulator's objective function associated with a (small) nonstochastic increase in the quality of the output. The right-hand side of that same expression is the associated marginal cost of such a nonstochastic increase in quality. For an interior solution to exist, the marginal utility of a nonstochastic increase in quality must be exactly equal to its cost. Hence, at the margin, there should be no potential social gains from a nonstochastic increase in product quality. This same condition also characterizes the first-best production of quality. It continues to apply in the second-best case because, there, nonstochastic changes in quality and payments can always be made incentive compatible. Hence, they are implementable provided they are consistent with individual rationality. As a result, in an interior equilibrium, no *sure* marginal social gains should ever be available at the optimal second-best contract. If they were, it would be optimal for the regulator to remove them by adjusting both the payment and the production nonstochastically.

This observation does not imply, however, that stochastic quality is at its first best level. Consider the second expression in (9). From IC_1, one sees that

$$y_2 - y_1 = \frac{\delta}{1 - \bar{\pi}_1^A} = \frac{z_2 - z_1}{1 - \bar{\pi}_1^A}.$$

Hence the right-hand side of the second expression in (9) measures the marginal cost to the regulator in state 2 income of maintaining incentive compatibility (preserving IC_1) when z_2 changes. The left-hand side of (9) represents the direct *ex post* benefit to the regulator from a change in quality in state 2. Thus, z_2 is chosen at the margin to balance the marginal cost of incentive compatibility against the regulator's gains from an increase in quality.

More precise evaluation of these first-order conditions requires an evaluation of P^o. Inspection of the regulator's concentrated objective function reveals that there are three possibilities; i) $P^o = \left\{\underline{\pi}_1^P\right\}$, ii) $P^o = \left\{\bar{\pi}_1^P\right\}$ and iii) $P^o = P$. These cases correspond to different rankings of $u(z_1^o)$ and $u(z_1^o + \delta^o) - \frac{\delta^o}{\left(1-\bar{\pi}_1^A\right)}$. When $u(z_1^o) > u(z_1^o + \delta^o) - \frac{\delta^o}{\left(1-\bar{\pi}_1^A\right)}$, $P^o = \left\{\underline{\pi}_1^P\right\}$. In contrast, when $u(z_1^o) < u(z_1^o + \delta^o) - \frac{\delta^o}{\left(1-\bar{\pi}_1^A\right)}$, the regulator uses the largest probability of state 1, $P^o = \left\{\bar{\pi}_1^P\right\}$. When $u(z_1^o) = u(z_1^o + \delta^o) - \frac{\delta^o}{(1-\bar{\pi}_1^A)}$, it follows that

$$u\left(z_2^o\right) - y_2^o = u\left(z_1^o\right) - y_1^o,$$

whence $P^o = \left[\underline{\pi}_1^P, \bar{\pi}_1^P\right]$. Just which of these three cases applies depends critically upon the parameters of the regulator's choice problem and upon the differences between the regulator's belief structure and that of the firm.

We now turn to the cases where (IC_2) and (IR) are binding at the optimum and where only (IR) is binding at the optimum. The technical analyses of these cases, which is discussed in more detail in the Appendix, largely parallel that of the previous case. We are able to establish (using a parallel notation) the following.

Proposition 6 *Suppose that $\alpha > \frac{\bar{\pi}_1^A}{(1-\bar{\pi}_1^A)}$ and (IC_2) and (IR) are binding at the optimum, then $z_2^o > z_1^o$ and a nonstochastic payment-quality (a quality standard) scheme is not optimal.*

Proposition 7 *Suppose that $\alpha > \frac{\bar{\pi}_1^A}{(1-\bar{\pi}_1^A)}$ and only (IR) is binding at the optimum. Then, a quality standard is not optimal and $z_2^o > z_1^o$ ($\delta^o > 0$).*

The latter result is also a consequence of condition $\alpha > \frac{\bar{\pi}_1^A}{(1-\bar{\pi}_1^A)}$. Suppose that the regulator implemented a quality standard in this case. But then because neither incentive compatibility constraint is binding, the regulator can increase her payoff marginally, without violating individual rationality or incentive compatibility, by modifying the original constant-payment contract so that payment and quality in state 1 are left unchanged while payment and quality in state 2 are increased by a sufficiently small amount. Thus, a quality standard is not optimal.

5 The Optimality of "Zero-risk" Quality Standards

"Zero-risk" quality standards can be a part of an optimal regulatory strategy under appropriate assumptions on the firm's belief structures. Taken together, Propositions 3, 5, 6, and 7 imply:

Proposition 8 *a) If* $\alpha > \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$*, then the second best quality vector satisfies* $z_2^o > z_1^o$*;*
b) If $\alpha < \frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}$*, then the second best quality vector satisfies* $z_2^o < z_1^o$*;*
c) If $\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A} \leq \alpha \leq \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$*, then the second best quality vector satisfies* $z_2^o = z_1^o$*.*

The intuition is relatively straightforward. If the technology is not inherently risky for the firm's prior (case c)), it is always less costly to produce the mean quality (as evaluated at that prior) with certainty than the stochastic quality. Thus, the regulator can replace the contract offered to the firm by one that involves the firm producing that mean quality with certainty while rewarding him or her with a nonstochastic payment. This is incentive compatible, and because cost has been lowered, it necessarily generates a surplus that the regulator can capture. In all other cases, however, there is always a cost incentive for the regulator to diverge from a "zero-risk" quality standard, which is always of first-order magnitude. Because a "zero-risk" quality standard is always incentive compatible, in the neighborhood of any such standard, incentive effects are negligible, and thus the first-order cost saving more than compensates for the incentive effects. Hence, the regulator optimally diverges from a "zero-risk" quality standard.

Thus, "zero-risk" quality standards can only be optimal when the marginal rate of transformation α belongs to a non-trivial set of relative probabilities $\left[\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}, \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}\right]$. As a point of comparison, notice that a quality standard is optimal in the first-best world only in the improbable cases characterized in Proposition 1. As noted earlier, the values of the parameters characterized in Proposition 1 have measure zero in the parameter space. Hence, in comparing the second-best scenario with the first-best, it seems that a "zero-risk" quality standard is more likely to emerge in the second-best case than in the first best case.

For purposes of comparison, it is interesting to contrast the results obtained here with those that would emerge for the more familiar case of expected-utility preferences. Expected-utility preferences, with a common belief structure, correspond to the special case of our model where there exist a common and singleton prior belief. That is, instead of a nontrivial set of relative probabilities, $\left[\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}, \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}\right]$, there exists a single unique relative probability common to both the regulator and the firm. In this polar case, "zero-risk" standards either never emerge or only emerge as optimal if the agent's technology is not inherently risky for the common prior shared by the principal and the agent (Quiggin and Chambers, 1998; Chambers and Quiggin, 2000). This is clearly an extreme, polar case. In fact, the set of cost parameters, α, for which this occurs also has measure zero in the parameter space. In the setting studied here, which while stylized seems more realistic, "zero-risk" quality standards optimally emerge if the firm's technology is not inherently risky for any of the priors in its zone of imprecision, $\left[\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}, \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}\right]$. So long as $\bar{\pi}_1^A > \underline{\pi}_1^A$, that is so long as there is ambiguity for the firm, the set of cost parameters, α, for which this will occur has strictly positive measure in the parameter space. Hence, in this sense, "zero-risk" quality standards are more likely to emerge in the presence of uncertainty or ambiguity than in its absence. Moreover, we can also conclude that the more imprecise the firm is about its ability to evaluate the likelihood of different states of the world occurring, the more likely it is that a standard will emerge. The link between the firm's (agent's) degree of imprecision is emphasized by noting that *when the firm's (agent's) beliefs are completely imprecise or ambiguous,* $\left[\underline{\pi}_1^A, \bar{\pi}_1^A\right] = [0,1]$, *a "zero-risk" quality standard is always an optimal policy.* In the expected-utility case, a "zero-risk" quality standard again occurs only for a set of measure zero in the parameter space.

The link between the firm's degree of imprecision and the optimality of "zero-risk" standards is particularly interesting. As we noted in our introduction, a traditional justification for a "zero-risk" approach is just such imprecision that arises from a lack of or fuzzy scientific knowledge, from poorly understood or unanticipated hazards, or from other forms of scientific uncertainty. Recall that in Main vs. Taylor (1986), the Supreme Court stated: "[The state] has a legitimate interest in guarding *against imperfectly understood environmental risks, despite the possibility that they may ultimately prove to be negligible*" (our italics). Even if one allows for the chance that a particular hazard might have negligible probability of occurring, our results establish that enough imprecision about that probability justifies a "zero-risk" posture as a constrained optimal response by the regulator.

6 Illustrative Example: Food-Safety Regulation

The bulk of current U.S. food safety regulation is carried out by the United States Department of Agriculture (USDA), overseeing meat, poultry and egg products, and the Food and Drug Administration (FDA),which is responsible for all other food products. In 1996, the Food Safety Inspection Service (FSIS) performed a major overhaul of US food-safety regulations by converting then existing command-and-control regulatory system to Pathogen Reduction/Hazard Analysis and Critical Control Point (HACCP). HACCP mandates that regulated companies form a set of rules to control food safety, establishes inspection policies and institutes testing requirements for *E. Coli* and Salmonella on meat and poultry (USDA). The latter are restrictions on the population density of colony forming units of the pathogen. A food producer is subject to penalties in case of failing to meet these testing requirements. Thus, our model where the transfer from the regulator to the firm depends on the observed quality of the good corresponds well with the existing regulatory system under HACCP of testing food products and penalizing poor quality.

The introduction of the HACCP regulations raised a plethora of new and difficult scientific and policy questions. One of the most important problems is the assessment of foodborne hazards associated with various strategies to control food safety. Difficulties with foodborne risk assessment arise due to imperfect hazard identification and inadequate scientific understanding of potential exposure and dose-response relationships (the link between levels of pathogens and the risk of various foodborne illnesses).[5] Put simply in our terminology, there is significant scientific uncertainty about foodborne hazards. FSIS and other government agencies have long recognized the existence of this uncertainty. For example, FSIS writes in its final rule on HACCP: "Perhaps of greater significance than the numerical estimate of risk is the uncertainty associated with the estimate. A fully developed risk characterization would include risk estimates and sensitivity/uncertainty analyses for alternative risk models and assumptions." (FSIS, 2006) The estimates of efficacy of various regulatory strategies

[5] FSIS (1996) points out these difficulties: "The present paucity of data to support a risk model for the major foodborne pathogens causing human disease limits the usefulness of quantitative risk assessment in the regulatory arena of meat and poultry inspection. It is unlikely that any single numerical constant will adequately describe the dose-response relationships for all pathogens associated with all of the products that FSIS regulates (CAST, 1994)."

on the level of foodborne pathogens and the costs of such strategies are also characterized by a substantial degree of uncertainty (see, e.g., Crutchfield et al., 1997). As a result, the existence of objective probabilities or the ability to determine informed subjective probabilities about foodborne hazards and the associated effects is at best problematic. Or, in other words, foodborne hazards often involve "considerable" Knightian uncertainty (Antle, 1996) which is the most important ingredient of our framework.

Thus, in terms of the formalisms of this paper, $\mathbf{z}$, could be identified with *a priori* uncertain food quality, $\mathbf{y}$, with the firm's *a priori* uncertain return for that quality, where actual or realized returns for the firm are dependent upon actual or realized food quality that is sold in the market after critical control point inspection. Quality standards, which can be identified with a command-and-control regulatory system as in existence before the introduction of the HACCP system, are then only optimal if the firm's marginal rate of transformation fell within $\left[\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}, \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}\right]$.

7 Relation to Previous Work

There is a small but growing literature on principal-agent relationships in the presence of uncertainty and ambiguity. Ghirardato (1994) examines optimal contracts when the principal and the agent have Choquet expected utility preferences (Schmeidler, 1989). In addition to proving the existence of solutions, Ghirardato (1994) investigates the validity of several properties of optimal contracts in a Grossman-Hart (1983) framework. Rigotti (2001) analyzes a principal-agent relationship in an environment where the agent has incomplete preferences as in Bewley (1986, 1987). Rigotti (2001) derives conditions under which an optimal incentive contract is *simple* in the sense that it consists of a base payment augmented by a bonus paid by the principal in some states of nature.

Mukerji (2003) studies optimal procurement contracts when contracting parties have $\alpha-$maximin expected utility preferences (Ghirardato, Maccheroni and Marinacci, 2004). For these preferences, beliefs are represented by a set of probability distributions and acts are evaluated by a convex combination of minimum and maximum expected utilities on that set of priors. The relative weight placed on the minimum expected utility is interpreted as the degree of ambiguity aversion. Mukerji (2003) shows that, provided the agent is more ambiguity averse than the principal, the power of the optimal procurement contract decreases as beliefs of the contracting parties become more ambiguous, where an increase in ambiguity of beliefs is modeled by expansion of the set of probabilities. Mukerji (2003) also demonstrates that the more ambiguity averse the agent relative to the principal the smaller is the power of the optimal incentive contract. Mukerji's (2003) results are directly relevant for the case at hand, and generally agree with our findings. "Zero-risk" quality standards, because they are characterized by fixed payments involve the lowest powered incentives possible.

Karni (2009) studies principal-agent relationships where the contracting parties have maximin expected utilities over action-dependent sets of priors. Karni (2009) also allows for outcome-dependent preferences over monetary payoffs. In addition to developing an axiomatization of the principal's and the agent's preferences, he explores some implications of ambiguity for design of optimal contracts in a framework with two actions and two possible outcomes.

While the motivation of this paper is very similar to those of the papers discussed above, the present work is different in several respects. Most importantly, we allow for greater variation in belief structures across the principal and the agent. In Ghirardato (1994) and Mukerji (2003) the principal and the agent have identical beliefs; in Rigotti (2001) the principal's beliefs are a subset of the agent's; Karni (2009) focuses on the two polar cases where either the principal's or the agent's set of beliefs is a singleton. In contrast, we do not place any prior restrictions on the principal's and the agent's beliefs because our primary goal is to investigate how differences in the parties' beliefs affect the optimal contract. Finally, these studies and our analysis are also different in terms of modeling certain aspects of the stochastic technology and attitudes towards risk and uncertainty.

8 Conclusion

This paper examines principal-agent relationships in uncertain environments where beliefs of the contracting parties are represented by sets of probabilities. In addition to characterizing the first-best and the second-best solutions, we have examined optimality of "zero-risk", fixed-payment schemes and the relationship between the first-best and the second-best solutions. In the second-best world, where the regulator can only contract on the quality of the good, a "zero-risk", fixed-payment scheme is optimal when the firm has beliefs that are imprecise enough that the firm's marginal rate of transformation in production belongs to the set of the firm's relative probabilities.

There are a few obvious extensions of our analysis. First is the extension to the general case with an arbitrary number of states. The second extension is to allow for a general state-contingent technology. One could also allow for general risk preferences on the firm's part. This latter generalization modifies our results minimally but leads to a considerably more involved exposition.

A particularly interesting, and at the same time challenging, extension of our model is the relaxation of the common-knowledge assumption (Aumann, 1976). Crew and Kleindorfer (page 70, 2006) call the assumption of common knowledge in a regulatory context its Achilles heel. In our context it would be interesting to explore how the results, especially optimality of the zero-risk regulation, would be affected if one weakened the assumption that the decision-makers' beliefs (their supports and/or likelihoods of different outcomes) are common knowledge. Because beliefs that are self-evident (for a definition, see, e.g., Geanakoplos, 1992) to both the regulator and the firm will be relatively imprecise, it seems that zero-risk regulation in such a setting would be more likely an optimal policy choice than under the corresponding scenario with common knowledge. To provide definitive answers to these questions one would need to explicitly model both a relaxation of the common-knowledge assumption and scientific uncertainty.[6]

[6] We would like to thank an anonymous Associate Editor for suggesting this discussion to us.

References

[1] Antle, J.M. (1996) "Efficient Food Safety Regulation in the Food Manufacturing Sector" *American Journal of Agricultural Economics* **78**, 1242-1247.

[2] Aumann, R. (1976) "Agreeing to Disagree" *The Annals of Statistics* **4**, 1236-39.

[3] Bewley, T.F. (1986) "Knightian Decision Theory: Part I" Cowles Foundation paper number 807.

[4] Bewley, T.F. (1987) "Knightian Decision Theory, Part II: Intertemporal Problems" Cowles Foundation paper numbers 835.

[5] Breyer, S. (1993) *Breaking the Vicious Circle: Toward Effective Risk Regulation.* Cambridge: Harvard University Press.

[6] Camerer, C.F. (1995) "Individual Decision Making" in J. Kagel and A. Roth (Eds.), *Handbook of Experimental Economics*, Princeton University Press, 587-703.

[7] Chambers, R.G. and J. Quiggin (2000) *Uncertainty, production, and choice, and agency: The state-contingent approach.* Cambridge; New York and Melbourne: Cambridge University Press.

[8] Chambers, R.G. and J. Quiggin (2005) "Incentives and Standards in Agency Contracts" *Journal of Public Economic Theory* **7**, 201-28.

[9] Chateauneuf, A., R.-A. Dana, and J.-M. Tallon (2000) "Optimal Risk-sharing Rules and Equilibria with Choquet Expected Utility" *Journal of Mathematical Economics* **34**, 191-214.

[10] Clarke, F.H. (1983) *Optimization and Nonsmooth Analysis.* Wiley, New York.

[11] Code of Federal Regulations (CFR), Title 21, available at http://www.access.gpo.gov/.

[12] Commission of the European Communities (CEC) (2000) *Communication on the Precautionary Principle.* Brussels, 02.02.2000.

[13] Crew, M.A. and P.R. Kleindorfer (2006) "Regulation, Pricing and Social Welfare" in M. A. Crew and D. Parker, eds., *International Handbook on Economic Regulation*. Edward Elgar Publishing, 63-81.

[14] Crutchfield, S.R., J.C. Buzby, T. Roberts, M. Ollinger, and C.-T. J. Lin (1997) "An Economic Assessment of Food Safety and Inspection: The New Approach to Meat and Poultry Inspection" Economic Research Service Agricultural Economic Report number 755.

[15] Geanakoplos, J. (1992) "Common Knowledge" *Journal of Economic Perspectives* **6**, 53-82.

[16] Food Safety Inspection Service, U.S. Department of Agriculture (1996) "The Final Rule on Pathogen Reduction and HACCP" *Federal Register* **61**, 38805-38855.

[17] Ghirardato, P. (1994) "Agency Theory with Non-Additive Uncertainty" Mimeo, California Institute of Technology.

[18] Ghirardato, P., F. Maccheroni and M. Marinacci (2004) "Differentiating Ambiguity and Ambiguity Attitude" *Journal of Economic Theory* **118**, 133-73.

[19] Gilboa, I. and D. Schmeidler (1989) "Maxmin Expected Utility with Non-unique Prior" *Journal of Mathematical Economics* **18**, 141-153.

[20] Grossman, S.J. and O.D. Hart (1983) "An Analysis of the Principal-agent Problem" *Econometrica* **51**, 7–45.

[21] Karni, E. (2009) "A Reformulation of the Maxmin Expected Utility Model with Application to Agency Theory" *Journal of Mathematical Economics* **45**, 97-112.

[22] Knight, F. (1921) *Risk, Uncertainty and Profit.* New York: Augustus M. Kelley.

[23] Lee, M. (2005) *EU Environmental Law: Challenges, Change and Decision-making.* Hart Publishing.

[24] Machina, M. and D., Schmeidler (1992) "A More Robust Definition of Subjective Probability" *Econometrica* **60**, 745-80.

[25] Mukerji, S. (2003) "Ambiguity Aversion and Cost-plus Procurement Contracts" Discussion paper number 171, Department of Economics, University of Oxford.

[26] Natural Resources Canada (NRC) (2000). available at http://www.nrcan-rncan.gc.ca/.

[27] Nichols, A.L. and R.J., Zeckhauser (1986) "The Perils of Prudence: How Conservative Risk Assessments Distort Regulation Regulation" *Regulation* **10**, 13-24.

[28] Quiggin, J. and R.G. Chambers (1998) "A State-Contingent Production Approach to Principal-Agent Problems with an Application to Point-Source Pollution Control" *Journal of Public Economics* **70**, 441-72.

[29] Racionero, M. and J. Quiggin (2006) "Fixed Wages and Bonuses in Agency Contracts: The Case of a Continuous State Space" *Journal of Public Economic Theory* **8**, 761-77.

[30] Rigotti, L. (2001) "Imprecise Beliefs in a Principal Agent Model" Mimeo, CentER, Tilburg University and Department of Economic, University of California, Berkeley.

[31] Rockafellar, R.T. (1970) *Convex Analysis.* Princeton University Press Princeton, N.J.

[32] Schmeidler, D. (1989) "Subjective Probability and Expected Utility without Additivity" *Econometrica* **57**, 571-587.

[33] Sunstein, C.R. (2007) *Worst-Case Scenarios.* Harvard University Press.

[34] Viscusi, W.K. and T. Gayer (2002) "Safety at Any Price?" *Regulation* **25**, 54-63.

[35] Vogel, D. (2003) "The Politics of Risk Regulation in Europe and the United States" *The Yearbook of European Environmental Law* **3**.

[36] Walley, P. (1991) *Statistical Reasoning with Imprecise Probabilities.* Chapman & Hall, London.

[37] Weisburger, J.H. (1996) "The 37 Year History of the Delaney Clause" *Experimental and Toxicologic Pathology* **48**, 183-188.

9 Appendix: Proofs of Main Results

Proof of Proposition 1 An interior production equilibrium requires that

$$\frac{u'(z_2^*)}{u'(z_1^*)} = \frac{\min\{\pi_1 : \boldsymbol{\pi} \in P^*\}}{\alpha \min\{\pi_2 : \boldsymbol{\pi} \in P^*\}}.$$

There are three relevant cases: $P^* = P$, $P^* = \left\{\left(\underline{\pi}_1^P, 1 - \underline{\pi}_1^P\right)\right\}$, and $P^* = \left\{\left(\bar{\pi}_1^P, 1 - \bar{\pi}_1^P\right)\right\}$, corresponding, respectively, to the cases where $u(z_2^*) - y_2^* = u(z_1^*) - y_1^*$, $u(z_2^*) - y_2^* < u(z_1^*) - y_1^*$, and $u(z_2^*) - y_2^* > u(z_1^*) - y_1^*$. We consider, without loss of generality, the first and the second. The third case is symmetric.

When $u(z_2^*) - y_2^* = u(z_1^*) - y_1^*$ and thus $u(z_2^*) - u(z_1^*) = y_2^* - y_1^*$, an interior production equilibrium requires that

$$\frac{u'(z_2^*)}{u'(z_1^*)} = \frac{\bar{\pi}_1^P}{\alpha\left(1 - \bar{\pi}_1^P\right)}.$$

Combined with the strict concavity of the regulator's utility structure this condition implies that $u(z_2^*) - u(z_1^*) > 0$ (whence $y_2^* - y_1^* > 0$) only if $\frac{\bar{\pi}_1^P}{\left(1-\bar{\pi}_1^P\right)} < \alpha$. By a parallel argument, it follows that $u(z_2^*) - u(z_1^*) = 0$ $(u(z_2^*) - u(z_1^*) < 0)$ only if $\frac{\bar{\pi}_1^P}{\left(1-\bar{\pi}_1^P\right)} = \alpha$ $\left(\frac{\bar{\pi}_1^P}{\left(1-\bar{\pi}_1^P\right)} > \alpha\right)$. When $u(z_2^*) - y_2^* < u(z_1^*) - y_1^*$, similar arguments establish that $u(z_2^*) - u(z_1^*) > 0$, $u(z_2^*) - u(z_1^*) = 0$, and $u(z_2^*) - u(z_1^*) < 0$ only if, respectively, $\frac{\underline{\pi}_1^P}{\left(1-\underline{\pi}_1^P\right)} < \alpha$, $\frac{\underline{\pi}_1^P}{\left(1-\underline{\pi}_1^P\right)} = \alpha$, and $\frac{\underline{\pi}_1^P}{\left(1-\underline{\pi}_1^P\right)} > \alpha$.

Proof of Proposition 2 It suffices to show that an interior equilibrium must involve $P^* \cap A^* \neq \varnothing$. From the text an interior payments equilibrium requires

$$\frac{\max\{\pi_1 : \boldsymbol{\pi} \in P^*\}}{1 - \min\{\pi_1 : \boldsymbol{\pi} \in P^*\}} = \frac{\min\{\pi_1 : \boldsymbol{\pi} \in A^*\}}{1 - \max\{\pi_1 : \boldsymbol{\pi} \in A^*\}}.$$

Suppose that $P^* \cap A^* = \varnothing$. There are two possible cases: $\max\{\pi_1 : \boldsymbol{\pi} \in P^*\} < \min\{\pi_1 : \boldsymbol{\pi} \in A^*\}$ or $\min\{\pi_1 : \boldsymbol{\pi} \in P^*\} > \max\{\pi_1 : \boldsymbol{\pi} \in A^*\}$. Consider the first; then the requirement for an interior payment equilibrium is given by

$$\max\{\pi_1 : \boldsymbol{\pi} \in P^*\} = \min\{\pi_1 : \boldsymbol{\pi} \in A^*\} \frac{1 - \min\{\pi_1 : \boldsymbol{\pi} \in P^*\}}{1 - \max\{\pi_1 : \boldsymbol{\pi} \in A^*\}},$$

whence

$$\frac{1-\min\{\pi_1:\boldsymbol{\pi}\in P^*\}}{1-\max\{\pi_1:\boldsymbol{\pi}\in A^*\}}<1,$$

which is a contradiction. Now consider the second; the requirement for an interior payment equilibrium is given by

$$1-\max\{\pi_1:\boldsymbol{\pi}\in A^*\}=(1-\min\{\pi_1:\boldsymbol{\pi}\in P^*\})\frac{\min\{\pi_1:\boldsymbol{\pi}\in A^*\}}{\max\{\pi_1:\boldsymbol{\pi}\in P^*\}},$$

whence,

$$\frac{\min\{\pi_1:\boldsymbol{\pi}\in A^*\}}{\max\{\pi_1:\boldsymbol{\pi}\in P^*\}}<1,$$

which is a contradiction.

Proof of Proposition 3 Combining (IC_1) and (IC_2) we obtain

$$\left(1-\bar{\pi}_1^A\right)(y_2-y_1)\geq z_2-z_1\geq\frac{\bar{\pi}_1^A}{\alpha}(y_2-y_1). \tag{10}$$

It follows immediately from (10) that the set of incentive compatible contracts that can implement (z_1,z_2) with $z_2>z_1$ is non-empty only if $\alpha\geq\frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$ $\left(\text{or, equivalently } \frac{\alpha}{1+\alpha}\geq\bar{\pi}_1^A\right)$. We now show, following Quiggin and Chambers (1998), that this same condition ensures that the set of incentive compatible contracts with $z_2>z_1$ is nonempty. From (10), condition $\frac{\alpha}{1+\alpha}\geq\bar{\pi}_1^A$ ensures that the first two incentive compatibility constraints are satisfied. To demonstrate that it also implies that the third is satisfied, multiply the first incentive compatibility constraint by $\left(1-\bar{\pi}_1^A\right)$ and the second by $\bar{\pi}_1^A$ and then sum the results to obtain:

$$\bar{\pi}_1^Ay_1+\left(1-\bar{\pi}_1^A\right)y_2-\alpha z_1-z_2\geq\left(1-\bar{\pi}_1^A\right)(y_1-\alpha z_1-z_1)+\bar{\pi}_1^A(y_2-\alpha z_2-z_2).$$

If $\frac{\alpha}{1+\alpha}\geq\bar{\pi}_1^A$, then

$$\left(1-\bar{\pi}_1^A\right)(y_1-\alpha z_1-z_1)+\bar{\pi}_1^A(y_2-\alpha z_2-z_2)\geq\bar{\pi}_1^Ay_2+\left(1-\bar{\pi}_1^A\right)y_1-\alpha z_2-z_1,$$

which verifies that all three constraints are satisfied.

Combining $(IC_1)'$ and $(IC_2)'$ we obtain

$$\frac{\underline{\pi}_1^A}{\alpha}\left(y_1 - y_2\right) \geq z_1 - z_2 \geq \left(1 - \underline{\pi}_1^A\right)\left(y_1 - y_2\right). \tag{11}$$

It also follows from (11) that the set of incentive compatible contracts that can implement (z_1, z_2) with $z_2 \leq z_1$ is non-empty if only if $\alpha \leq \frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A}$. Thus, using (10) and (11), one sees that when $\frac{\underline{\pi}_1^A}{1-\underline{\pi}_1^A} \leq \alpha \leq \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$, only production vectors (z_1, z_2) with $z_2 = z_1$ are implementable.

Proof of Proposition 5 Since the case where $z_1^o = 0$ is trivial we prove the result only for $z_1^o > 0$. From (8) we have that

$$V'\left(z_1^o, \delta^o = 0; \hat{z}_1 = 1, \hat{\delta} = 0\right) = \left[u'(z_1^o)\right] - (1 + \alpha) = 0 \tag{12}$$

for a solution with $z_1^o > 0$ and $\delta^o = 0$. (8) also implies that $\delta^o = 0$ if and only if

$$V'\left(z_1^o, \delta^o = 0; \hat{z}_1 = 0, \hat{\delta} = 1\right) = \min\left\{(1 - \pi_1)\left[u'(z_1^o) - \frac{1}{(1 - \bar{\pi}_1^A)}\right] : \pi_1 \in P^o\right\} \leq 0,$$

where z_1^o is implicitly given by (12). Hence, $\delta^o = 0$ if and only if $\alpha \leq \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$. Recall that we have supposed that $\alpha > \frac{\bar{\pi}_1^A}{\left(1-\bar{\pi}_1^A\right)}$. Hence, $\delta^o > 0$.

10 Appendix: Alternative Binding Constraints

10.1 (IC_2) and (IR) are binding at the optimum

Because the derivations for this case are very similar to the previous case, we do not provide as many details. In this case, the payment schemes that are implementable for (z_1, z_2) are given by:

$$\begin{aligned}
\bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 &= y_2 - \alpha z_2 - z_2, \\
\bar{\pi}_1^A y_1 + \left(1 - \bar{\pi}_1^A\right) y_2 - \alpha z_1 - z_2 &= \bar{u}.
\end{aligned}$$

Solving gives:

$$\begin{aligned}
y_2 &= \bar{u} + (1 + \alpha) z_2, \\
y_1 &= \bar{u} + (1 + \alpha) z_2 - \frac{\alpha \delta}{\bar{\pi}_1^A}.
\end{aligned}$$

where $\delta \equiv z_2 - z_1$. By substituting these expressions into the regulator's objective function we obtain a concentrated objective function that can be used to characterize the equilibrium contract structure by the optimal choice of z_1 and δ with $\delta \geq 0$

$$V(z_1, \delta) = \min \left\{ \pi_1 \left[u(z_1) + \frac{\alpha \delta}{\bar{\pi}_1^A} \right] + (1 - \pi_1) u(z_1 + \delta) : \pi_1 \in \left[\underline{\pi}_1^P, \bar{\pi}_1^P \right] \right\} - (\bar{u} + (1 + \alpha)(z_1 + \delta)),$$

This concentrated objective function is concave in (z_1, δ). Hence, any local solution must also be a global solution to this problem. Denote:

$$(z_1^o, \delta^o) \in \underset{(z_1, \delta)}{\arg\max} V(z_1, \delta).$$

and, with a slight abuse of notation,

$$P^o \equiv \arg\min \left\{ \pi_1 \left[u(z_1^o) + \frac{\alpha \delta^o}{\bar{\pi}_1^A} \right] + (1 - \pi_1) u(z_1^o + \delta^o) : \pi_1 \in \left[\underline{\pi}_1^P, \bar{\pi}_1^P \right] \right\}.$$

Optimality requires that the directional derivative of $V(z_1,\delta)$, whose existence is guaranteed by concavity and denoted by $V'\left(z_1,\delta;\hat{z}_1,\hat{\delta}\right)$, be non-positive for all possible directions $\left(\hat{z}_1,\hat{\delta}\right)$. Calculation shows that

$$V'\left(z_1^o,\delta^o;\hat{z}_1,\hat{\delta}\right) = \min\left\{\pi_1\left[u'(z_1^o)\hat{z}_1+\frac{\alpha\hat{\delta}}{\bar{\pi}_1^A}\right]+(1-\pi_1)u'\left(z_1^o+\delta^o\right)\left(\hat{z}_1+\hat{\delta}\right):\pi_1\in P^o\right\} \tag{13}$$

$$-(1+\alpha)\left(\hat{z}_1+\hat{\delta}\right)\leq 0 \text{ for all } \left(\hat{z}_1,\hat{\delta}\right).$$

Using condition (13) and arguments almost identical to those used in obtaining Proposition 5, leads to Proposition 6.

Thus, here too, the optimal contract structure is characterized by a stochastic payment scheme $(y_2^o > y_1^o)$ and a stochastic quality vector $(z_2^o > z_1^o)$. As in the previous case, inspection of the regulator's concentrated objective function reveals that there are three possibilities; i) $P^o = \left\{\underline{\pi}_1^P\right\}$, ii) $P^o = \left\{\bar{\pi}_1^P\right\}$ and iii) $P^o = \left[\underline{\pi}_1^P,\bar{\pi}_1^P\right]$. These cases correspond to different rankings of $u(z_1^o)+\frac{\alpha\delta^o}{\bar{\pi}_1^A}$ and $u(z_1^o+\delta^o)$.

10.2 Only (IR) is binding at the optimum

In this case the only binding constraint is

$$\bar{\pi}_1^A y_1 + \left(1-\bar{\pi}_1^A\right)y_2 - \alpha z_1 - z_2 = \bar{u}.$$

The analysis in this case is similar to the derivations for the first-best problem and is, therefore, omitted.

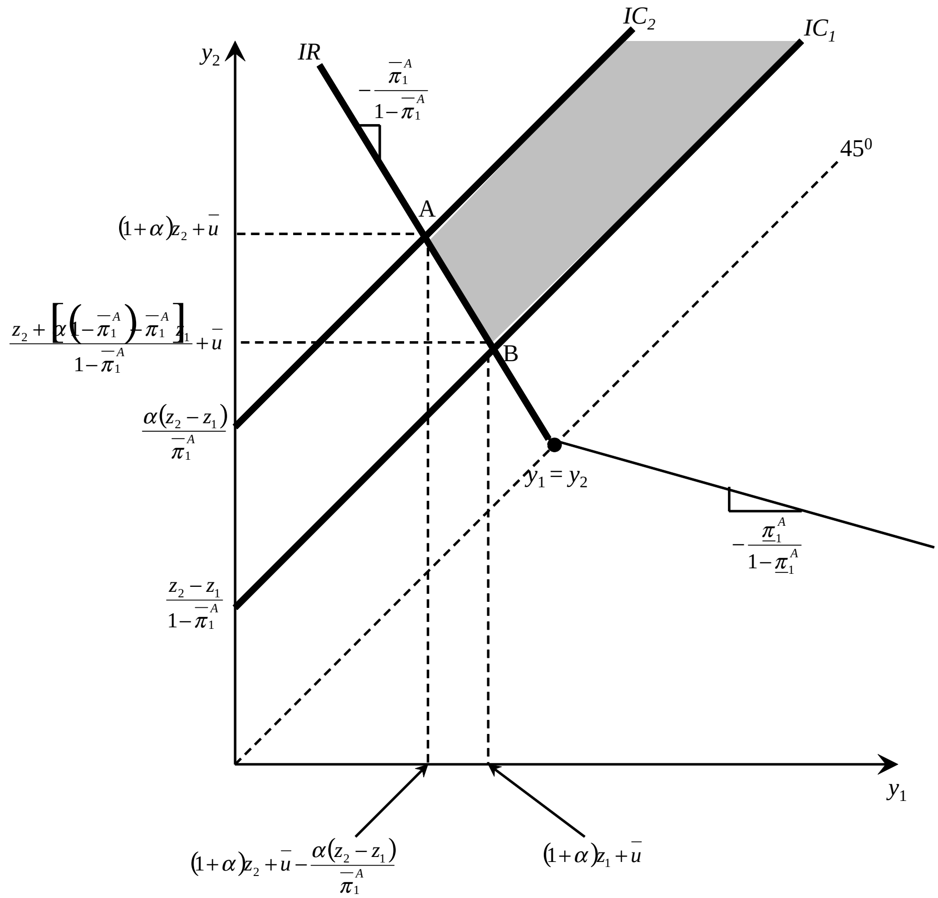

Figure 1. Individual Rationality and Incentive Compatibility Constraints.