Risk Sharing in Coasian Contracts[1]

Arthur A. Small
Columbia University

Joshua Zivin
Columbia University

March 28, 2000

Abstract

The Coase Theorem is analyzed in a setting in which pollution damages are a stochastic function of emissions and of natural environmental variability (e.g., weather). When pollution damages are stochastic, emissions create financial risks. Pollution levels allowed under Coasian contracts then in general depend on agents' risk appetites, and on the initial configuration of property rights and bargaining power. In this case, resource allocation decisions are not separable from the legal institutions that allocate risks, nor from the financial institutions that facilitate risk transfer. In particular, improvements in institutions for environmental risk management generally induce greater levels of pollution. (JEL Q2, K1, D8)

The Coase Theorem can be viewed as a statement that technical externalities can be internalized through a clear assignment of property rights that supports private bargaining between parties (Coase 1960, Stigler 1966). Under a costlessly functional system of private rights in pollution, parties will always negotiate contracts that specify the production of that unique level of pollution that maximizes aggregate wealth; the assignment of rights merely affects how this wealth is divided between agents. The result has received a great deal of attention in the literatures on environmental economics, law and economics, and legal scholarship generally. Subsequent research has clarified the conditions needed to support the Coase result. The theorem does not generally hold, for example, when agents' preferences exhibit wealth effects (Hurwicz 1995), and when contracting parties possess asymmetric information (Farrell 1987).

With rare exception, discussion of the Coase Theorem and its implications have assumed a world of certainty. In many cases, however, the damages caused by pollution or other externalities depend on the realization of exogenous stochastic variables. Of these, among the most important govern variation in the natural environment. A chemical discharge from a factory may befoul a waterway, but the damage it imposes on downstream users will also depend on the volume of streamflow, a function of recent rainfall. Nitrous oxides emitted from a power plant may contribute to the formation of ground-level urban ozone — but the rate of formation depends also on current temperature. Natural catastrophes such as hurricanes and earthquakes can turn routine production into sources of major pollution problems. In 1999, for example, Hurricane Floyd's visit to North Carolina sent thousands of hog carcasses, and millions of gallons of animal waste, into the streams and onto the beaches and even the front porches of that state.

When pollution damages depend on exogenous stochastic variables, Coasian contracts that affect production and pollution levels also affect the production of risk. The effect of risk on private Coasian contracting has received surprisingly little attention in the economics literature. One exception is the work of Greenwood and Ingene (1978), apparently the first to consider Coasian contracting between risk-averse agents when realized damages are stochastic in pollution levels. They showed that when an externality creates risk for the impacted party,

emissions will in general be lower when the externality is internalized through merger, than when it is internalized through a contract specifying allowed pollution levels and payment flows. The Greenwood-Ingene framework is quite limited, however. In effect, the authors provide an existence proof, an example of a single case in which risky liabilities lead to the contravention of the Coase result. They do not develop a framework that allows examination of how equilibrium pollution levels depend on agents' risk appetites, how these change as property rights and bargaining power are reallocated between agents, and on how contracts will be structured to allow for the sharing of risk.

In this paper, we address three questions. First, how does the presence of risk affect the design of Coasian contracts — in particular, equilibrium pollution levels? Second, how will agents use contingent contracts to share risks, and what are the effects if they cannot? Third, we examine how equilibrium contracts are affected by the availability of financial institutions that facilitate risk transfer to third parties. These questions are analyzed with the aid of a model of Coasian contracting between risk averse agents in a stochastic environment. Our framework allows for an identification of how equilibrium pollution levels depend on the allocation of property rights and bargaining power, and on the risk preferences of agents.

Several insights emerge from the analysis. First, we find that under very general conditions the Coase Theorem does not hold if damages are stochastic in emissions and if agents are risk averse. This result is closely related to, but distinct from, the role of wealth effects in the deterministic setting. Second, we show how the parties to the contract will share risks if they are able, and what happens if they cannot. When contingent contracting is infeasible, the pollution level specified in the equilibrium contract depends on the risk attitude of the risk-bearing agent. This level is insensitive to the initial allocation of rights only in the case for which the risk-bearing agent exhibits constant absolute risk aversion (CARA). For the arguably more realistic case of decreasing absolute risk aversion (DARA), equilibrium pollution levels increase as the risk-bearing agent extracts more of the surplus created by the contract. When contingent contracting is feasible, risk can be shared but not eliminated. Thus, the sensitivity of equilibrium pollution levels to the assignment of property rights and bargaining power remains. Indeed, the Coase theorem holds in this setting if and only

if agents share the same coefficient of absolute risk aversion. More generally, equilibrium pollution levels depend on agents' ability to extract surplus and their relative willingness to bear risk. Finally, we show that in essentially all cases, improvements in the financial markets that allow for management of environmental risk through third parties will increase pollution levels in equilibrium.

The balance of the paper is organized as follows. Section 1 presents a model of Coasian contracting between risk averse agents when pollution damages are a stochastic function of emissions levels. In section 2 we derive equilibrium contracts for the case in which realized damages cannot be directly monitored, and in the case when realized damages can be used as the basis for a damage-contingent contract. Section 3 analyzes how the form of equilibrium contracts is affected by the introduction of environmental securities (i.e. securities with payoffs contingent on the realization of stochastic environmental events). Section 4 concludes.

1   The Model

There are two agents, a polluter and a victim. The polluter chooses output level x, causing the victim dollar damages in the amount $(1+\theta)x$, where θ is a random variable representing the effect of environmental stochasticity on damages, drawn from a distribution with mean $\bar{\theta}$ and variance $\sigma_\theta^2 > 0$. By rescaling x as necessary, we may assume without loss of generality $\bar{\theta} = 0$. Aggregate damages therefore have mean x and variance $\sigma_\theta^2 x^2$.[1] Production at level x creates a benefit to the polluter in the amount $\pi(x)$, where π is differentiable, concave and single-peaked, with $\pi(0) = 0$ and $\pi'(0) > 1 + \sigma_\theta$.[2] Absent contracting, polluter profits are maximized when $x = \bar{x}$.

---

[1]This framework comprehends as cases both those of accidental and planned damages. Suppose, for example, that with probability $1/K$ a rare environmental event (e.g., an earthquake) will cause K dollars of damage per unit production level x, but that damages will be zero otherwise. This situation could be represented by letting θ have a binary distribution such that $\theta = K - 1$ with probability $1/K$ and $\theta = -1$ with probability $(K - 1)/K$. A situation in which damages are continuous in emissions (e.g., urban ozone formation as a function of both NO_x emissions and temperature) could be represented by letting θ have a continuous distribution.

[2]More generally, we may think of x as encapsulating both the scale of the firm's activity and (negatively) the effort the firm deploys toward pollution abatement or accident avoidance. The only real restrictions are that the profit function be concave in x, and that $\pi'(0)$ be sufficiently large to rule out shut-downs in equilibrium. This last assumption avoids Starrett-type non-convexities in the polluter's utility function (Starrett 1972).

Polluter and victim may elect to enter into a contract specifying an allowed level of pollution x. In general, contracts may call for a combination of lump-sum and damage-contingent payments. Let $\tau(\theta)$ denote the net payment from victim to polluter, as a function of realized environmental shocks. (If property rights lay initially with the victim, then $\tau(\theta) \leq 0$.) For tractability, we restrict our attention to cases in which payments are linear in damage. Let β denote the share of variable damages retained by the victim under the contract. Then τ takes the form $\tau(\theta) = \gamma - (1 - \beta)\theta x$, where γ represents a lump-sum payment, and $1 - \beta$ governs the level of risk borne by the polluter.

A contract comprises a triple (x, γ, β). Under the contract, the victim's wealth W_v takes the form $W_v = W_0 - \gamma - x - \beta\theta x$, where W_0 is the victim's wealth initially, i.e. in the absence of pollution. We assume $W_0 > (1 + \sigma_\theta)\bar{x}$. The victim's stochastic wealth therefore has mean $\bar{W}_v = W_0 - \gamma - x$ and standard deviation $\sigma_v = \beta\sigma_\theta x$. Likewise, polluter wealth is given by $W_p = \pi(x) + \gamma - (1 - \beta)\theta x$, with mean $\bar{W}_p = \pi(x) + \gamma$ and standard deviation $\sigma_p = (1 - \beta)\sigma_\theta x$.

The preferences of polluter and victim over stochastic wealth levels are represented by a Mean-Standard Deviation (MS) utility function. For $i = p, v$, agent i's preferences over stochastic wealth are ranked according to the function $U_i(\bar{W}_i, \sigma_i) = \bar{W}_i^{\alpha_i} - \sigma_i^{\alpha_i}$.[3] The parameter α_i is an index of absolute risk aversion. The cases $\alpha > 1$, $\alpha = 1$, and $\alpha < 1$ correspond to decreasing, constant, and increasing absolute risk aversion. The restriction that the exponent on mean wealth equals that on variance of wealth corresponds to an assumption that both agents' preferences display constant relative risk aversion. Our analysis will concentrate on the cases for which $\alpha_i \geq 1$ for $i = p, v$ (i.e., the realistic cases involving DARA or CARA preferences).

The assumption $W_0 > (1 + \sigma_\theta)\bar{x}$ on the victim's initial wealth implies that all contracts must satisfy a participation constraint $U_v > 0$. Likewise, we assume that the polluter always has the option to cease production. Hence all contracts must also satisfy a participation constraint on polluter utility $U_p \geq 0$.

---

[3]Since wealth for both polluter and victim are linear in θ for all contracts (x, γ, β), preference ordering under $U(\bar{W}, \sigma)$ maximization is consistent with preference ordering under expected utility maximization. See Meyer (1987, pg. 422) and Saha (1997, pg. 772).

2 Analysis

The terms of the contract that will obtain in equilibrium can in principle depend on the initial assignment of property rights. If the space of possible contracts are viewed as points in an Edgeworth Box with dimensions x, γ, and β, an assignment of property rights to the victim corresponds to a no-contract default at the point $(0, 0, 0)$; an assignment of rights to the polluter corresponds to a default of $(\bar{x}, 0, 0)$.

An initial allocation of property rights determines the set of feasible contracts. The contract that obtains in equilibrium will depend both on this initial allocation, and on agents' relative bargaining power, their ability to extract the surplus created through the contract. An equilibrium contract that delivers utility u_v to the victim can equivalently be derived as the solution to the problem of maximizing the polluter's utility, such that victim utility is at least u_v:[4]

$$\max_{x,\gamma,\beta} \bar{W}_p^{\alpha_p} - \sigma_p^{\alpha_p} \tag{1}$$

$$\text{s.t. } \bar{W}_v^{\alpha_v} - \sigma_v^{\alpha_v} \geq u_v \tag{2}$$

Thus, a given reservation utility level u_v encapsulates both the initial allocation of rights and each agent's bargaining power. Varying u_v traces out the entire contract curve.

The contract curve contains, therefore, all equilibrium contracts that can be obtained under any initial allocation of rights, given any distribution of bargaining power. Rather than specify a bargaining game and solving for a single equilibrium contract, our strategy will be to derive expressions for the entire contract curve, i.e. the set of all Pareto-optimal contracts.

2.1 No Contingent Contracting

Ideally, a contract will specify a level of liability contingent on actual, realized damages. A contingent contract provides a mechanism by which the two parties can share the risk associated with stochastic damages. This arrangement will in general allow for efficient risk-bearing. However, in some situations, technical or legal constraints may preclude the use of

---

[4]Naturally, the same contract curve can also be derived by solving the dual problem in which victim utility is maximized subject to a constraint on polluter's reservation utility.

contingent contracts. For some toxic-related health risks (for example, when toxic exposure increases the incidence of cancer), the victim may not even know directly the causal link between contamination and subsequent adverse health outcome. In other cases, it may be costly or impossible to measure the environmental factors that govern realized damage with sufficient credibility to provide the basis for a contract. Cases in which contingent contracts are infeasible arise, therefore, in a range of real-world problems. Furthermore, the analysis of these cases provide a simplified setting in which the model can be examined for insights that are portable to the more general case.

We start by examining the case in which the polluter bears all risk. The equilibrium contract (x, γ) can be found by maximizing (1) subject to (2), and the additional constraint $\beta = 0$. The problem yields the following first-order condition and binding constraint for an interior solution:

$$\pi'(x) = 1 + A_p \sigma_\theta \tag{3}$$

$$\bar{W}_v^{\alpha_v} = u_v \ . \tag{4}$$

Here, $A_i \equiv \left(\sigma_i/\bar{W}_i\right)^{\alpha_i - 1}$ measures agent i's risk attitude, the marginal rate of substitution between increases in mean wealth and decreases in standard deviation of wealth, i.e. the slope of the indifference locus in W_i-σ_i space (Saha 1997).

Equation (3) describes equilibrium pollution levels. It states that the benefit created by the marginal unit of pollution must just offset the expected damage caused, plus a premium that depends on the associated increase in aggregate risk, weighted by the polluter's willingness to absorb that risk. Since the polluter has the option to shut down if his utility $u_p = \bar{W}_p^{\alpha_p} - \sigma_p^{\alpha_p}$ would otherwise be negative, we are assured that $\sigma_p/\bar{W}_p \leq 1$. Hence (3) implies that for an interior solution $1 \leq \pi'(x) \leq 1 + \sigma_\theta$. Since π is assumed to be concave, the condition $\pi'(0) > 1 + \sigma_\theta$ is sufficient to insure that pollution levels are positive in equilibrium. Equation (4) simply restates the constraint on victim utility for the case in which the victim bears no risk.

In this setting, the Coase Theorem corresponds to a condition that equilibrium pollution levels are constant along the contract curve. By totally differentiating the system (3—4)

with respect to endogenous variables (x, γ) and exogenous parameters $(\sigma_\theta, \alpha_p, \alpha_v, W_0, u_v)$, we derive an expression for dx/du_v, which describes how a change in the victim's reservation utility affects equilibrium pollution levels:

$$\frac{dx}{du_v} = (\alpha_p - 1)\frac{1}{\bar{W}_p} \cdot \frac{1}{u_v} / |H| \tag{5}$$

where $|H|$, the determinant of the bordered Hessian matrix, is negative, ensuring an interior maximum (see Appendix).

The Coase Theorem corresponds to the condition that $dx/du_v = 0$ along the contract curve. When the polluter exhibits decreasing absolute risk aversion ($\alpha_p > 1$) this equality, which we term the Coase Condition, will never hold. Indeed, the Coase Condition obtains if and only if the polluter exhibits constant absolute risk aversion, $\alpha_p = 1$.

Similar logic applies to the case where the victim bears all risk, i.e. when the constraint $\beta = 1$ is imposed. It is straightforward to show that in this case the rate at which equilibrium pollution levels change with respect to changes in the victim's reservation utility can be expressed as follows:

$$\frac{dx}{du_v} = -(\alpha_v - 1)\frac{1}{\bar{\bar{W}}_v}\left(\frac{1}{\sigma_v^{\alpha_v} + u_v}\right) / |H| \tag{6}$$

where again $|H|$, the determinant of the bordered Hessian matrix, is negative (see Appendix). In this case, the Coase Condition will hold if and only if the victim exhibits constant absolute risk aversion, $\alpha_v = 1$. These results are summarized in the following proposition.

Proposition 1 *In this framework, if the victim bears all risk and has DARA preferences, then equilibrium pollution levels are an increasing function of the victim's reservation utility. If the polluter bears all risk and has DARA preferences, then equilibrium pollution levels are a decreasing function of the victim's reservation utility.*

In short, when damages are a stochastic function of pollution and the contracting parties cannot insure each other against the associated risks, then equilibrium pollution levels depend on the risk-appetite of the risk-bearing agent. Unless this agent has CARA preferences, his risk-appetite depends on his wealth level, which depends in turn on the initial allocation of

pollution rights, and on his bargaining power. If the victim cannot shed risk via a contingent contract, then improving victim utility requires either a decrease in the lump-sum payment delivered to the polluter, or a decrease in the pollution level. The latter both improves mean wealth and decreases the variance of wealth. Therefore, the trade-off between x and γ will depend on the victim's appetite for risk. If the victim has DARA preferences, then as his wealth increases, he will be more willing to accept the risk of higher levels of pollution in return for lower lump-sum payments. Likewise, if the polluter must bear all risk, she makes a similar risk-reward trade-off. As the victim's wealth increases along the contract curve, the polluter's wealth decreases. If she has DARA preferences, she then becomes less willing to accept risk, so decreases pollution levels.

A refinement of the Coase Theorem to the case of uncertain damages follows immediately from (5) and (6).

Proposition 2 *In this framework, when agents cannot strike contingent contracts, the Coase Theorem holds if and only if the risk-bearing agent exhibits constant absolute risk aversion.*

2.2   Contingent Contracting

If realized damages $(1 + \theta)x$ can be monitored directly and are verifiable, then the parties will in general be able to strike a contingent contract. A contingent contract allows both agents to improve their outcomes by sharing the aggregate pollution risk.

To derive the equilibrium contract in this expanded space, we solve (1) subject to (2). This procedure yields first-order conditions for an interior solution:

$$\pi'(x) = 1 + A_p \sigma_\theta \tag{7}$$

$$A_p = A_v \tag{8}$$

and the binding constraint on victim utility,

$$\bar{W}_v^{\alpha_v} - \sigma_v^{\alpha_v} = u_v \ . \tag{9}$$

Equation (7), describing equilibrium pollution levels, is unchanged from (3) above. The new condition (8) describes an optimal risk sharing rule: the mix of certain and damage-

contingent payments should be allocated to the two agents so as to equalize their risk attitudes. In the special case in which both agents have the same coefficient of absolute risk aversion ($\alpha_p = \alpha_v$), (8) implies that $\beta/(1-\beta) = \bar{W}_v/\bar{W}_p$: the ratio of the agents' risk shares equals the ratio of their respective mean wealth levels. Likewise, $\beta/(1-\beta) < \bar{W}_v/\bar{W}_p$ when $\alpha_p > \alpha_v > 1$ and $\beta/(1-\beta) > \bar{W}_v/\bar{W}_p$ when $\alpha_v > \alpha_p > 1$.

In this setting, the Coase Theorem corresponds to a condition that equilibrium pollution levels are constant along the contract curve. By totally differentiating the system (7—9) with respect to endogenous variables (x, γ, β) and exogenous parameters $(\sigma_\theta, \alpha_p, \alpha_v, W_0, u_v)$, we derive an expression for dx/du_v, which describes how a change in the victim's reservation utility affects equilibrium pollution levels:

$$\frac{dx}{du_v} = (\alpha_p - 1)(\alpha_v - 1)z \left(\frac{1}{(1-\beta)\bar{W}_v} - \frac{1}{\beta\bar{W}_p} \right) / |H| \tag{10}$$

where $z \equiv \alpha_v \sigma_v^{\alpha_v-1}/(\sigma_v^{\alpha_v} + u_v)$ and where $|H|$, the determinant of the bordered Hessian matrix, is positive, ensuring an interior maximum (see Appendix).

When both agents exhibit decreasing absolute risk aversion ($\alpha_p > 1$ and $\alpha_v > 1$), the Coase Condition holds if and only if $\beta/(1-\beta) = \bar{W}_v/\bar{W}_p$. For the case $\alpha_p = \alpha_v$, this equality is equivalent to first-order condition (8), and so holds automatically. For $\alpha_p \neq \alpha_v$, (7) and the Coase Condition together imply $\bar{W}_p - \sigma_p = \bar{W}_v - \sigma_v = 0$. These equalities imply in turn that $u_p = u_v = 0$; this condition cannot hold along an entire interval of points on the contract curve.

Our results can be summarized in the following proposition.

Proposition 3 *Let (x, γ, β) denote the pollution level, lump-sum transfer, and risk-sharing rule, respectively, for the equilibrium contract arising in the model described above. If both agents' preferences exhibit decreasing absolute risk aversion ($\alpha_p > 1$ and $\alpha_v > 1$), and if $\pi'(0) > 1 + \sigma_\theta$, then the following hold:*

1. *Equilibrium pollution levels are positive ($x > 0$), and bounded by the condition $1 < \pi'(x) < 1 + \sigma_\theta$.*

2. *Equilibrium contracts always involve risk sharing ($0 < \beta < 1$).*

3. *An increase in the victim's reservation utility induces an increase (a decrease, no change) in equilibrium pollution when the victim's coefficient of absolute risk aversion is greater than (less than, equal to) that of the polluter: $sgn(dx/du_v) = sgn(\alpha_v - \alpha_p)$.*

Intuitively, one can view 1 and $1 + \sigma_\theta$ as lower and upper bounds, respectively, on the aggregate perceived damage caused by each unit of pollution, taking into account agents' aversion to risk. Claim 1 then says simply that the aggregate benefit of the marginal unit of pollution must lie within these bounds. Claim 2 says that contingent contracts, when feasible, are always preferred to lump-sum contracts, as they allow pollution-related risks to be spread across both agents. To understand Claim 3, it helps to think of the reservation utility that a contract must deliver to the victim as a form of wealth. As the victim's reservation utility increases, this wealth shifts from the polluter to the victim. Since both agents' have DARA utility, this shift raises the victim's willingness to shoulder risk, while lowering that of the polluter. The net impact on pollution levels depends on how these two effects combine to alter the "aggregate" risk appetite of the parties. When $\alpha_v > \alpha_p$, the risk tolerance of the victim is growing faster than that of the polluter is shrinking; the net effect is an increase in aggregate risk tolerance. When $\alpha_v < \alpha_p$, the net effect has the opposite sign.

Proof. We first rule out equilibria in which production ceases, by showing that there would always exist a Pareto-superior contract. Consider the case in which the victim has the property rights initially. If $x = 0$, then $\pi(x) = 0$ and $\sigma_p = 0$; the polluter's participation constraint $\bar{W}_p \geq \sigma_p$ then implies $\gamma = 0$. Under this contract, $U_p = 0$ and $U_v = W_0^{\alpha_v}$. Suppose now that instead $x = \varepsilon$, a small positive number, so that $\pi(x) \simeq \pi'(0)\varepsilon > (1 + \sigma_\theta)\varepsilon$. Let $\beta = 1$, and let $\gamma = -(1 + \sigma_\theta)\varepsilon$. Under this alternative contract, $U_p = (\pi(\varepsilon) + \gamma)^{\alpha_p} > 0$, and $U_v = (W_0 + \varepsilon\sigma_\theta)^{\alpha_v} - (\varepsilon\sigma_\theta)^{\alpha_v} > W_0^{\alpha_v}$.[5] A similar proof rules out equilibria with $x = 0$ for the case in which the polluter has the initial allocation of property rights.

Since by assumption $W_0 > (1 + \sigma_\theta)\bar{x}$, the victim's participation constraint $U_v = \bar{W}_v^{\alpha_v} -$

---

[5]In general, for $\alpha > 1$, $a > 0$ and $b > 0$ we have that $(a + b)^\alpha > a^\alpha + b^\alpha$. To see this, suppose (without loss of generality) that $a \geq b$, and let $\delta = a/b$. Since the function $f(y) = y^\alpha$ is strictly convex for $\alpha > 1$, a Taylor series expansion shows that $(\delta + 1)^\alpha > \delta^\alpha + \alpha\delta^{\alpha-1} > \delta^\alpha + 1$, where the last inequality holds since $\delta \geq 1$. Multiplying this inequality through by b^α yields the claim.

$\sigma_v^{\alpha_v} \geq (W_0 - \bar{x})^{\alpha_v} - (\bar{x}\sigma_\theta)^{\alpha_v} > 0$ insures that $0 \leq \sigma_v/\bar{W}_v < 1$. Since $\alpha_p > 1$ and $\alpha_v > 1$, (8) then implies $0 \leq A_v = A_p < 1$. Claim 2 then follows since $x > 0$: if to the contrary $\beta = 0$ then $A_v = 0 < A_p$, while $\beta = 1$ would imply $A_p = 0 < A_v$.[6] Hence $0 < A_v = A_p < 1$. Claim 1 then follows immediately from (7). Finally, under the condition that $\alpha_p > 1$ and $\alpha_v > 1$, equation (10) shows that $sgn(dx/du_v) = sgn\left(\beta/(1-\beta) - \bar{W}_v/\bar{W}_p\right) = sgn(\alpha_v - \alpha_p)$, proving Claim 3. ∎

It turns out that the assumption of decreasing absolute risk aversion is essential to the results of Proposition 3. If the polluter, say, has CARA utility ($\alpha_p = 1$), while the victim has DARA utility ($\alpha_v > 1$), then in equilibrium we must have $A_v < A_p = 1$. Under any arrangement, the victim is strictly more willing to bear risk at the margin than is the polluter. The option to strike a contingent contract will not be exercised: the victim will in equilibrium bear all the risk ($\beta = 1$), and pollution levels will be given by the condition $\pi'(x) = 1 + A_v\sigma_\theta$. The case in which $\alpha_p > 1$ and $\alpha_v = 1$ is symmetric. Finally, if both agents have CARA utility, then $A_v = A_p = 1$: since both agents share an identical willingness to bear risk that is unaffected by wealth levels, the exact mix of certain and contingent payments is indeterminate. All we can say in this case is that, given a reservation utility u_v for the victim, γ and β will follow the linear relationship $W_0 - \gamma - x - \beta x\sigma_\theta = u_v$, where equilibrium pollution x is a constant determined by the condition $\pi'(x) = 1 + \sigma_\theta$.

Drawing on this discussion, the last claim in Proposition 3 leads to a generalized statement of the Coase Theorem in stochastic environments under contingent contracting.

Corollary 4 *In this framework, if both agents' preferences exhibit constant or decreasing absolute risk aversion and if linear contingent contracts are feasible, then the Coase Theorem holds if and only if both agents have the same coefficient of absolute risk aversion.*

---

[6] A contract with $\beta < 0$ is Pareto-dominated by one with $\beta = 0$; the move decreases both agents' risk without sacrifice of mean return. Equilibria involving $\beta > 1$ are similarly ruled out.

3   Effect of Environmental Securities on Equilibrium Contracts

In many realistic cases, agents' opportunities for risk management are not confined to bilateral contracts between affected agents. To the extent that agents can use financial products to shift their risks to others, their willingness to accept risky contracts increases. A full examination of how risk affects Coasian contracting must attend, therefore, to the opportunities the financial system provides for the management of risky liabilities.

Insurance policies are standard tools for managing liability risk. Suppose that Brown, who operates a hog farm adjacent a river, and Charlotte, a lovely city downstream, have entered into a contract (either voluntary or mandated by law) that calls for Brown to compensate a negotiated share of Charlotte's damages in case of a manure spill. If a hurricane sends Brown's hog manure into Charlotte's drinking water, part or all of Brown's liability may be covered by general business liability insurance. More generally, two parties to a Coasian contract may settle on a risk-sharing rule, then contract with third-party insurers to off-load part or all of their remaining risk.

As an instrument for managing environmental risks, insurance carries a significant advantage of "targeting": payouts are a function of actual realized damages (in our model, θx). Further, because insurance premiums will in general be an increasing function of the risk taken on (x), insurance contracts provide incentives to internalize externalities.[7] However, the standard problems of moral hazard, adverse selection, transaction costs, *ex post* opportunism and other market imperfections will in general imply that full insurance is unavailable at actuarially fair rates.

An additional problem will frequently arise when the risks arising from environmental variation are positively correlated across contracts. Indeed, this situation is likely to be the norm for many environmental risks: if many agents in a region are damaged by a hurricane or frost, they probably suffer in tandem. To the degree that environmental risks are systematic rather than idiosyncratic, the pooling of risks through an insurance company intermediary

---

[7]For an informal discussion of how environmental liabilities can be managed through insurance, see Freeman and Kunreuther (1997).

will be ineffective.

Risks arising from stochastic variation in the natural environment can also be managed, in some cases, through the use of environmental derivatives. For example, Farmer Brown could buy securities that pay off if and only if a hurricane above a certain severity makes landfall in North Carolina during a specified time period. To the extent that these securities' payoffs are correlated with his own damages, he becomes partly insured against his contractual risk. (On such uses of environmental securities, see Chichilnisky and Heal (1998) and the articles in Froot (1999) and in Himick and Bouriaux (1998).) More generally, improvements in the financial system that increase agents' ability to hedge environmental risks will increase the willingness of risk-averse agents to accept more such risk. This development will, it is claimed, lead to higher levels of pollution in equilibrium.

To demonstrate the claim formally, let Ω denote the space of financial portfolios that can be formed from the set of securities available in the market, i.e. the closure of that set under linear combinations. We assume away complications involving imperfections in financial markets: there are no arbitrage opportunities, transaction costs, borrowing constraints or limits on short sales. We also assume the existence of a risk-free security $1 \in \Omega$, with unit payoff in every state of the world. The problem confronting each agent is one of choosing a portfolio $\omega \in \Omega$ that provides an optimal hedge against the agent's residual risk. An ideal hedge would be a portfolio the payoff on which is perfectly correlated with the environmental shock θ, and which is available at zero (or even negative) cost. In general, with incomplete markets, the choice of an optimal hedge will depend on each portfolio's price, expected return, variance of return, and the covariance of that return with the environmental shock θ, as well as on the agent's appetite for risk.

Deriving closed-form analytic solutions for this general problem is not always possible, however. To simplify matters, we impose the additional restriction that the environmental shock θ be uncorrelated with aggregate market returns.[8] This assumption is reasonable

---

[8]Even in the general case, in which the environmental risk covaries with the market's return, we can proceed by decomposing θ into a systematic component θ_m and an idiosyncratic component θ_0, such that the latter is uncorrelated with aggregate market returns. Our immediate problem is that of choosing an optimal hedge against idiosyncratic risks associated with realizations of θ_0. This problem is separable from that of choosing an optimal market portfolio, given security prices and returns, and agent attitudes toward

for most environmental risks. Environmental variation that exacerbates or mitigates the damages caused by a given level of contamination are generally local or regional in scope. Local environmental variables such as precipitation, temperature, wind speed and direction, or tidal action will generally have no impact on aggregate global market returns. It is true that large-scale natural perturbations such as hurricanes, earthquakes, tornados, serious frosts, or shifts in the ENSO (El Niño—Southern Oscillation) cycle may have significant economic impacts in certain regions or industries. Even these shocks, however, will generally have only negligible effects on the performance of a globally diversified financial portfolio.[9]

Formally, let $m \in \Omega$ denote a pricing kernel for Ω: for all $\omega \in \Omega$, the cost $c(\omega)$ of purchasing the portfolio ω is given by $c(\omega) = E[m\omega]$. The assumption that the environmental shock is uncorrelated with market returns can be restated as the claim that $Cov(m, \theta) = E[m\theta] - E[m]E[\theta] = 0$. Since $E[\theta] = 0$ by construction, we also have $E[m\theta] = 0$.

To find the optimal hedging portfolio within Ω, let $\tilde{\theta}$ denote a projection of θ onto Ω under the inner product function $\langle \omega_1, \omega_2 \rangle = E[\omega_1 \cdot \omega_2]$. That is, θ is decomposed as $\theta = \tilde{\theta} + \varepsilon$, where $\tilde{\theta} \in \Omega$, $E[\varepsilon\omega] = 0$ for all $\omega \in \Omega$, and $\tilde{\theta}$ is chosen to minimize $Var(\varepsilon)$. Clearly, $\tilde{\theta}$ is a zero-cost portfolio: $c(\tilde{\theta}) = E[m\tilde{\theta}] = E[m\theta] - E[m\varepsilon] = 0$, where the last equality holds since, by construction, $E[\omega\varepsilon] = 0$ for all $\omega \in \Omega$. Also, $\tilde{\theta}$ has zero expected return: $E[\tilde{\theta}] = E[\theta] - E[1 \cdot \varepsilon] = 0$, where the last equality holds since $1 \in \Omega$. Let $\sigma_{\tilde{\theta}}^2 = Var(\tilde{\theta})$, and let $\rho = Cov(\theta, \tilde{\theta})/\sigma_\theta \sigma_{\tilde{\theta}}$ denote the coefficient of variation of $\tilde{\theta}$ with respect to θ.

Suppose the polluter in our model has struck a contract (x, γ, β) with the victim that calls for her to assume some pollution related financial risk (i.e. for which $\beta < 1$). She can then reduce her residual exposure by adding copies of $\tilde{\theta}$ to her portfolio. Since $c(\tilde{\theta}) = 0$, she may acquire any number $\tilde{x}_0$ of these securities without changing her expected wealth. Rescale units by letting $\tilde{x} = \tilde{x}_0(\sigma_{\tilde{\theta}}/\sigma_\theta)$, so that $Var(\tilde{x}_0\tilde{\theta}) = \tilde{x}^2\sigma_\theta^2$. Her random wealth then becomes $\tilde{W}_p = \pi(x) + \gamma - (1 - \beta)\theta x + \tilde{\theta}\tilde{x}(\sigma_\theta/\sigma_{\tilde{\theta}})$, with mean $\bar{W}_p = \pi(x) + \gamma$ and variance

holding non-diversifiable risk. However, the extent of the systematic risk θ_m will influence how the agent assembles the rest of her market portfolio and, therefore, her willingness to shoulder pollution-related risks.

[9]The claim cannot be stated categorically, however. It is certainly conceivable that an earthquake unfolding directly under Silicon Valley that measured 9.0 on the Richter Scale could induce a sharp drop in NASDAQ share prices. Indeed, some observers cite the Kobe Earthquake as a significant factor contributing to Japan's continued economic difficulties. Furthermore, Hurricane Andrew of 1994, the most expensive natural disaster in U.S. history, did place substantial stress on the global reinsurance industry.

$\tilde{\sigma}_p^2 = [(1-\beta)^2 x^2 + \tilde{x}^2 - 2(1-\beta)x\tilde{x}\rho]\,\sigma_\theta^2$. Since the securities have no effect on the polluter's mean wealth, she will choose $\tilde{x}$ to minimize the (disutility of) variation in wealth: $\partial\tilde{\sigma}_p^2/\partial\tilde{x} = 0$ implies $\tilde{x} = \rho(1-\beta)x$. By acquiring $\tilde{x}_0 = (1-\beta)x \cdot Cov(\theta,\tilde{\theta})/Var(\tilde{\theta})$ copies of the security $\tilde{\theta}$, the polluter reduces the variance in her wealth from σ_p^2 to $\tilde{\sigma}_p^2 = \sigma_p^2\,(1-\rho^2)$. By the same argument, the victim will purchase $\beta x \cdot Cov(\theta,\tilde{\theta})/Var(\tilde{\theta})$ of the securities, yielding wealth with variance $\tilde{\sigma}_v^2 = \sigma_v^2\,(1-\rho^2)$.

Here, ρ is a summary measure of how well the set of available securities can be combined to approximate the distribution of damages caused by variation in the natural environment. As ρ ranges from 0 to 1, the financial system goes from being useless, to providing perfect hedging against the environmental risk. Developments in the financial system can have real effects on equilibrium contracts. Assuming that each agent hedges his residual risk by forming portfolios as described above, emissions x will be determined by the equilibrium condition $\pi'(x) = 1 + \tilde{A}_i\sigma_\theta$, where i is the index of a risk-bearing agent and where $\tilde{A}_i = A_i\,(1-\rho^2)^{(\alpha_i-1)/2}$ is the new risk attitude of this agent. As the financial system improves, both agents become more accepting of risk. Since π' is decreasing in x, equilibrium pollution levels x are increasing in ρ. In the limiting case in which the financial system offers agents the opportunity to hedge their risks completely (as ρ grows toward unity), both agents will behave as if they are risk neutral. The classical Coase result is restored.

The introduction of financial hedging will also usually alter the contractual allocation of risks. In particular, when the contract calls for damage-contingent payments, the new risk sharing rule $\tilde{A}_v = \tilde{A}_p$ will in general identify a different equilibrium value of β, except in the case where agents share the same coefficient of absolute risk aversion ($\alpha_v = \alpha_p$).

4 Conclusion

In most situations, the damages caused by pollution depend stochastic environmental factors. Variability in temperature, rainfall, wind, and other weather patterns can have profound effects on the costs generated by externalities. Since many (arguably, most) economic agents are averse to risk, a true measure of "pollution damages" must, therefore, include any disutility associated with volatility in the natural environment. In this paper, we analyze how

risk averse agents will use bilateral Coasian contracts to manage such uncertain externalities. When information problems make it costly to verify realized damages, agents will be confined to contracts specifying lump-sum payments in exchange for changes in the allowed level of output. In these cases, one agent retains all risk. The risk appetite of this agent will govern equilibrium outcomes. When this agent's risk tolerance varies with wealth, equilibrium pollution levels will depend on initial property allocations. The ability to strike contingent contracts allows agents to share, but not eliminate, pollution-related risks. Contingent contracts, when feasible, will be designed both to select an optimal pollution level and to spread optimally the resulting financial risk. The level of pollution that arises will depend on agents' risk appetites, will in general be lower than that arising when damages are certain, and will in most cases increase as property rights and bargaining power are shifted toward the agent with a greater coefficient of absolute risk aversion. In either case, the Coase Theorem typically does not hold.

The central issue for law and policy is how these risk effects shape Coasian contracts in practice. To address this question, we must attend to the opportunities agents have for managing their risks through financial markets. Environmental securities, including such novel instruments as hurricane catastrophe bonds and degree-day derivatives, have recently emerged as a means to hedge risks generated by stochasticity in the natural world. As the menu of environmental financial products expands, agents will increasingly be able to shed some of their pollution-related risks onto capital markets. From this vantage point, each agent's exposure to pollution-related risks can be viewed as part of his overall portfolio of assets and liabilities. The introduction of environmental hedges pushes out the mean-variance frontier of the agent's portfolio space: he responds by accepting more pollution risk. In short, improvements in financial markets can lead to higher levels of pollution.

The analysis has important implications. First, when realized pollution damages depend on stochastic environmental outcomes, legal rules and institutions for allocating liability have real effects on pollution levels, even in the absence of any barriers to contracting. If we care about pollution for reasons beyond those of the bargaining agents, the assignment of property rights will impact social welfare. Hence, the design of legal rules should not be

independent of polluter and victim risk attitudes. Second, environmental outcomes are not separable from the financial system: the effect of a liability rule depends on the state of financial markets. Thus, legal rules may need to evolve in response to an enrichment in the menu of financial mechanisms for managing environmental risks.

A Appendix

Here we confirm the claims that the first order conditions derived in the main body of the text are sufficient to insure unique equilibria.

A.1 Case: No Contingent Contracting, Polluter Bears All Risk ($\beta = 0$)

A sufficient condition for the existence of a unique maximum of the objective function for this case is that $|H|$, the determinant of the Bordered Hessian matrix, must be negative. The condition in this case is

$$
\begin{aligned}
|H| &= \left| \begin{array}{cc} \frac{-\pi''(x)}{\pi'(x)-1} + (\alpha_p - 1)\left[\frac{1}{x} - \frac{\pi'(x)}{W_p}\right] & -\frac{\alpha_p - 1}{W_p} \\ -\frac{\alpha_v}{W_v} & -\frac{\alpha_v}{W_v} \end{array} \right| \\
&= \frac{\alpha_v}{\bar{W}_v}\left\{ \frac{\pi''(x)}{\pi'(x)-1} - (\alpha_p - 1)\left[\frac{1}{x} - \frac{\pi'(x)-1}{W_p}\right] \right\} < 0.
\end{aligned}
$$

Since $\pi'(x) - 1 \leq \sigma_\theta$, $\pi''(x) < 0$, and $\sigma_p = x\sigma_\theta$, the condition $|H| < 0$ holds if $\bar{W}_p \geq \sigma_p$. This last inequality simply restates the polluter's participation constraint $U_p \geq 0$. ∎

A.2 Case: No Contingent Contracting, Victim Bears All Risk ($\beta = 1$)

Again, a unique maximum of the objective function is assured if $|H|$, the determinant of the Bordered Hessian matrix, is negative. For this case, the condition is that

$$
\begin{aligned}
|H| &= \left| \begin{array}{cc} \frac{\pi''(x)}{1-\pi'(x)} + (\alpha_v - 1)\left[\frac{1}{x} - \frac{1}{W_v}\right] & -\frac{\alpha_v - 1}{W_v} \\ -\alpha_v/\bar{W}_v - z\sigma_\theta & -\alpha_v/\bar{W}_v \end{array} \right| \\
&= -\frac{\alpha_v}{\bar{W}_v}\left\{ \frac{\pi''(x)}{1-\pi'(x)} + (\alpha_v - 1)\frac{1}{x} \right\} - (\alpha_v - 1)\frac{z\sigma_\theta}{\bar{W}_v} < 0
\end{aligned}
$$

where $z \equiv \alpha_v \sigma_v^{\alpha_v - 1}/(\sigma_v^{\alpha_v} + u_v) > 0$. Given that $\pi''(x) < 0$ and $\pi'(x) > 1$, this inequality holds by inspection. ∎

A.3 Case: Contingent Contracting $(0 < \beta < 1)$

In this case, the Bordered Hessian is a 3×3 matrix. We must show that its determinant $|H|$ is positive. This expression,

$$
|H| = \begin{vmatrix}
\frac{\pi''(x)}{\pi'(x)-1} + (\alpha_p - 1)\left[\frac{\pi'(x)}{\bar{W}_p} - \frac{1}{x}\right] & \frac{\alpha_p - 1}{\bar{W}_p} & \frac{\alpha_p - 1}{1-\beta} \\
(\alpha_p - 1)\left[\frac{1}{x} - \frac{\pi'(x)}{\bar{W}_p}\right] - (\alpha_v - 1)\left[\frac{1}{x} + \frac{1}{\bar{W}_v}\right] & -\frac{\alpha_p - 1}{\bar{W}_p} - \frac{\alpha_v - 1}{\bar{W}_v} & -\frac{\alpha_p - 1}{1-\beta} - \frac{\alpha_v - 1}{\beta} \\
-\frac{\alpha_v}{\bar{W}_v} - z\sigma_\theta\beta & -\frac{\alpha_v}{\bar{W}_v} & -z\sigma_\theta x
\end{vmatrix},
$$

can after considerable algebraic manipulation be expressed as the sum of three terms: $|H| = H_1 + H_2 + H_3$. A sufficient condition for $|H| > 0$ is for all three terms to be positive.

The first term takes the form

$$
H_1 = \frac{\pi''(x)}{\pi'(x) - 1} \frac{\alpha_v}{\bar{W}_v} \left\{ (\pi'(x) - 1)\, x \left[\frac{\alpha_p - 1}{\bar{W}_p} + \frac{\alpha_v - 1}{\bar{W}_v} \right] - \left[\frac{\alpha_p - 1}{1 - \beta} + \frac{\alpha_v - 1}{\beta} \right] \right\}.
$$

Recall that $\pi'(x) - 1 \le \sigma_\theta$, $\beta x \sigma_\theta = \sigma_v$, and $(1-\beta)x\sigma_\theta = \sigma_p$. Denoting the expression in braces by $\{\cdot\}_1$, we have $\{\cdot\}_1 \le x\sigma_\theta(\alpha_p - 1)\left(1/\bar{W}_p - 1/\sigma_p\right) + x\sigma_\theta(\alpha_v - 1)\left(1/\bar{W}_v - 1/\sigma_v\right)$. This expression is negative since $\bar{W}_p \ge \sigma_p$ and $\bar{W}_v \ge \sigma_v$ by the agents' participation constraints. Since also $\pi''(x) < 0$, $H_1 > 0$.

To show that the second term,

$$
H_2 \equiv (\alpha_p - 1)(\alpha_v - 1)\frac{\alpha_v}{\bar{W}_v}\left(\bar{W}_v + (\pi'(x) - 1)\beta x\right)\left\{\bar{W}_p - \pi'(x)x\right\},
$$

is positive, it suffices to show that the expression in braces, $\{\cdot\}_2 \equiv \bar{W}_p - \pi'(x)x = \pi(x) + \gamma - \pi'(x)x$, is positive. Recall that $x \ge 0$. At $x = 0$, $\{\cdot\}_2 = \gamma$, which is non-negative by the polluter's participation constraint. Furthermore, $\{\cdot\}_2$ is increasing in x: $\partial\{\cdot\}_2/\partial x = \pi'(x) - \pi'(x) - \pi''(x)x > 0$. Hence $\{\cdot\}_2$, and H_2, are positive.

To show that the last term,

$$
H_3 \equiv (\alpha_p - 1)(\alpha_v - 1)\frac{\alpha_v}{\bar{W}_v}\left[\left(\frac{1}{(1-\beta)\,x} + \frac{1}{\beta\bar{W}_p}\right)\left\{1 - \frac{\beta x\,(\pi'(x) - 1)}{\bar{W}_v}\right\}\right],
$$

is positive, it suffices to show that the expression in braces, $\{\cdot\}_3 \equiv 1 - \beta x\,(\pi'(x) - 1)/\bar{W}_v$, is positive. We know that $\pi'(x) - 1 \le \sigma_\theta$ and that $\beta x\sigma_\theta = \sigma_v$. Thus $\{\cdot\}_3 > 1 - \sigma_v/\bar{W}_v > 0$ by the victim's participation constraint. Therefore $H_3 > 0$, which completes the proof. ∎

References

[1] G. Chichilnisky and G. Heal. Managing unknown risks: The future of global reinsurance. *Journal of Portfolio Management*, 24(4), 1998.

[2] R. H. Coase. The problem of social cost. *Journal of Law and Economics*, 3(October):1–44, 1960.

[3] J. Farrell. Information and the coase theorem. *Economic Perspectives*, 1(2, Fall):113–129, 1987.

[4] P. K. Freeman and H. Kunreuther. *Managing Environmental Risk Through Insurance.* Studies in Risk and Uncertainty. Kluwer Academic Publishers, Boston, 1997.

[5] K. Froot, editor. *The Financing of Catastrophic Risk.* The University of Chicago Press, Chicago and London, 1999.

[6] P. H. Greenwood and C. A. Ingene. Uncertain externalities, liability rules, and resource allocation. *American Economic Review*, 68(3):300–310, 1978.

[7] M. Himick and S. Bouriaux, editors. *Securitized Insurance Risk: Strategic Opportunities for Insurers and Investors.* Glenlake Publishing Co., Chicago, 1998.

[8] L. Hurwicz. What is the coase theorem? *Japan and the World Economy*, 7(1):49–74, 1995.

[9] J. Meyer. Two-moment decision models and expected utility maximization. *American Economic Review*, 77(3):421–430, 1987.

[10] A. Saha. Risk preference estimation in the nonlinear mean standard deviation approach. *Economic Inquiry*, XXXV(October):770–782, 1997.

[11] D. Starrett. Fundamental nonconvexities in the theory of externalities. *Journal of Economic Theory*, 4:542–552, 1972.

[12] G. J. Stigler. *The Theory of Price.* Macmillan, New York, 3rd edition, 1966.