Accidents Happen:

The Effect of Uncertainty on Environmental Policy Design

Thomas Sproul[1] and David Zilberman [2]

Dept. of Agricultural & Resource Economics

UC Berkeley

Selected Paper prepared for presentation at the Agricultural & Applied Economics Association's 2011 AAEA & NAREA Joint Annual Meeting, Pittsburgh, Pennsylvania, July 24-26, 2011

---

[1] Corresponding author. Email: sproul@berkeley.edu.

[2] The authors would like to thank Gordon Rausser, David Sunding, Ben Hermalin, Peter Berck, Sofia Villas-Boas, Jenny Ifft, Leslie Martin, Joanne Lee, Frankie Le, Kyriakos Drivas, Gal Hochman, Howard Chong, David Roland-Holst and Joachim Otte for their patient reading, insightful comments, and feedback. We would also like to thank the members of the ARE Department Seminar, Pizza and Policy, and Zilberman's Graduate Research Seminar. Research leading up to this article was funded in part by a grant from UN-FAO and UK-DFID.

Abstract

Major externality cases are random accidents, which are not adequately addressed by the deterministic environmental policy literature - that of Pigouvian taxes, abatement subsidies and cap-and-trade. We consider a risk-neutral industry where firms control the probability and severity of accidents by preventive and responsive choices, but asymmetric information means government only observes outcomes. We show that even without intervention, some care will be taken, however - we identify three policies that lead to the optimal solution: strict liability, a stochastic subsidy, and a mandatory mutual insurance scheme. The subsidy policy may be very costly to taxpayers, especially when prevention affects the probability of accident occurrence, and strict liability may be excessively draconian; polluters are also victims and liabilities must exist regardless of adherence to professional standards of care. Thus, we propose a revenue-neutral liability-pooling scheme which plays a similar role to tradable pollution permits in the deterministic case.

Introduction: The Accidents Problem

Accidents happen to all of us. In the course of our everyday lives, or in the course of operating a business, we all bear the risk of things going wrong. An **externality** is created when accidents harm others, or impose costs on them, and one does not account for the social costs of risky behavior - as when preventive activities may reduce the likelihood or severity of harms. The element of randomness inherent in **stochastic externalities**, which include environmental accidents such as the BP spill in the Gulf of Mexico, the chemical plant explosion at Bhopal, or livestock disease outbreaks of avian or swine flu, has important implications for environmental policy design.

Classical environmental economic policy was developed around externalities created with relative certainty, like the emissions of a coal plant, which are necessary to energy production at some level. We focus on the ‘holy trinity’ of environmental economic policy: Pigouvian taxes, abatement subsidies, and systems of **tradable pollution permits** (TPPs), as described in Baumol and Oates (1988). While all of these policies can achieve optimal levels of pollution, the difference lies in the allocation of resources; taxes take money out of the polluting industry, subsidies transfer money in, and tradable permits (or ‘cap-and-trade’) keep all monies within the industry, but shift funds from ‘dirtier’ to ‘cleaner’ firms in equilibrium. Due to the differing resource allocations among optimal policies, political economic considerations may affect which policy is actually used, with industry preferring subsidies, then cap and trade, then taxes (Buchanan and Tullock, 1975).

In this paper, we introduce a conceptual framework for regulating an industry of risk neutral, stochastically polluting firms. We develop a general model allowing for information asymmetry, prevention and cleanup actions (both *ex ante* and *ex post*), and for the exact prevention mechanism to affect the probability and/or severity of accidents. The information asymmetry exists because polluters have better information; it may be very costly to know the costs of prevention and cleanup technologies for every firm in a large industry, and pollution being stochastic means the regulator cannot necessarily infer the agents’ activities. Thus, we focus on optimizing price-based policies of taxes and subsidies because quantity-based regulations, such as TPPs, standards or command-and-control policies, are not always possible (Zivin et al., 2005). In this setup, we demonstrate uncertainty-adjusted versions of tax and subsidy policies, but show that uncertainty may exacerbate discrepancies between resource allocations when compared to the deterministic setting.

The key results show that tax policy must take the form of **strict liability**, equivalent to a pricing system where the accidental polluter owes for the social cost of all pollution, even though they do not supply the random trigger. In essence, agents' actions are key components of the **risk-generation process** (Lichtenberg and Zilberman, 1988), a concept used to model the development and interaction of risk factors over time. The agents play a key - and exclusive - role in influencing if, and how much, damage will occur, even though they are also accident victims themselves. Strict liability can be a politically (or some argue, morally) unpalatable solution, because accidental polluters need not have violated social norms or professional standards of care when they are forced to pay.

While tax policies take money out of the polluting industry (funds which might theoretically be used to compensate the people harmed), subsidy policies establish incentives by transferring money into the industry; this relationship holds true for established, deterministic policies, as well as for their stochastic counterparts which we introduce here. We demonstrate that subsidy transfers into the industry may increase under uncertainty because subsidies must pay for reductions below a threshold *in every period*, even when no accident occurs. Given the possibility of ability-to-pay or political economic constraints for the optimal tax and subsidy policies, we introduce a novel policy of mutual insurance, which attains optimal outcomes and leads to a resource allocation similar to that of tradable permits.

Literature Background

This paper is, by no means, the first to examine regulation of stochastic externalities. However, we introduce a key consideration that is not present in the relevant literature: the tradeoff between care and cleanup. The basic idea here is that economic agents (also individuals, firms, or farms, herein) face choices of *ex ante* care (also caution, precaution, prevention) to decrease the probability of an accident and/or its severity, and *ex post* containment (also cleanup, mitigation, abatement) to lessen the harm to others once an accident occurs. Accidents releasing hazardous materials into the environment follow precisely this pattern; precautionary activities may include the imposition of safety controls and worker training, while mitigation efforts include notifying the public, or proper authorities, and actual cleanup of pollutants released.

Existing literature relevant to the subject can be grouped into three basic categories, beyond the classical externalities literature derived from deterministic frameworks. First, there is a literature of optimal second-best regulation under uncertainty, where the uncertainty arises from measurement error or information asymmetry, starting with

Weitzman (1974) and including the agency literature of Holmstrom (1982) and many more. These articles do not focus on externalities *per se*, but instead deal with regulatory information problems, whereby efficiency conditions are difficult to meet. Our framework incorporates some of these informational challenges to inject realism into the analysis; we allow agents' actions and their **types**, or inherent characteristics, to be unobservable to the regulator - and we demonstrate conditions under which the policies we propose are robust to the resulting problems of **moral hazard** (actions unobservable) and **adverse selection** (types unobservable). We also show that policies must remain adaptive to **state-dependent** outcomes, where the optimal levels of pollution and containment efforts will depend on the state of nature, because price or quantity controls fixed *ex ante* are constrained to be second-best (Weitzman).

Beyond issues of uncertain information, there is literature on regulating stochastic externalities which is concerned with uncertainty due to randomness of outcomes. However, these papers often use simplified models, where accidents have fixed severity and where the key results are driven by risk preferences. For example, Just and Zilberman (1979) rely on risk preference to demonstrate the asymmetry of taxes and subsidies for stochastic externalities, as do Zivin et al. (2005) when evaluating Coasean bargaining under a continuum of property rights regimes. While it is clear that the Coase Theorem can hold for stochastic externalities under risk neutrality, the implications for classical environmental policy and the associated political economic considerations have not been explored. Risk neutrality is an appropriate starting point for the analysis because our framework focuses on large, infrequent environmental accidents, of which the perpetrators are generally large corporations (who act as profit maximizers). The lack of risk preference is also a staple of welfare economics; it implies that ‘pure transfers’ do not change aggregate welfare - which, in this context, allows for a straightforward comparison of stochastic externality policies vis-a-vis deterministic ones.[3]

Finally, there is the classic accidents literature,[4] which has evaluated legal liability standards according to economic efficiency measures, like Pareto optimality. However, the frameworks have focused only on *ex ante* care choices (Edlin (1994), Shavell (1985) and (1987), Landes and Posner (1983)),[5] or *ex post* remediation of damages (Polinsky and Shavell (1994), Innes (1999b)), which limits the scope of their policy prescriptions. For example, the idea of **due care**, where liability is limited or zero when precautions conform

---

[3] Zivin and Small (2003) have shown that actions can be invariant under risk-aversion for both parties, though it requires strong assumptions: all agents have known, identical, constant absolute risk aversion utility functions. Outside this special case, it is clear that analytical results will be driven by assumptions about utility functions and endowments.

[4] e.g., Shavell (1980), Polinksy (1980), Cooter and Porat (2000), among many.

[5] These papers do, however, focus on the role of outside parties (victims) in the probability of an accident occurring. Here, we focus on one party’s ability to control the extent of damages incurred by others.

to industry or social norms, focuses on cases where inputs (risky behaviors) are observable.

Unfortunately, the due care approach can only be optimal when *ex post* containment efforts are nonexistent, or regulated similarly. Our paper innovates by introducing a general, and complex, risk generation function which allows for *ex ante* and *ex post* damage control activities, and for prevention that acts on probability or on severity of accidents. We focus on liability-based regulation to address cases where outcomes are observable, but risky inputs may not be. This approach is especially useful for heterogeneous polluting firms - because varying risk profiles imply varying standards of due care, with attendant information requirements that may be costly to obtain.

We start the analysis by developing optimal stochastic externality taxes, to show that only a system with the marginal incentives of strict liability can obtain socially optimal outcomes when standards or quantity-based policies are infeasible. That is, the regulator needs no information about agents' choices, as long as the damages can be measured and traced to their source.[6] Unfortunately, strict liability can mean imposing a large fine on a

---

[6] For a thorough review of optimal regulation when damages cannot be traced, see Segerson (1988), Swierzbinski (2002), and Millock et al (2002) for theory on collective punishment, Ribaudo and Caswell (1999) for documentation of actual policy using the threat thereof, and Hamilton and Zilberman (2006) for an evaluation of voluntary traceability to capture consumer willingness-to-pay.

party that has already suffered a substantial loss. This may be socially distasteful, politically infeasible, or practically impossible if there are bankruptcies, which is the so-called **judgment-proof** problem (see for example, Shavell (1986), Beard (1990), Polborn (1998), Innes (1999a)).

When a strict liability regime is infeasible, we derive a subsidy policy which can attain the social optimum even when agents must voluntarily opt-in, or when they are judgment-proof. To demonstrate this policy in a broad setting, we provide a general framework for the accident mechanism and the effects of precautionary activity, agnostically allowing for precaution to affect the yes/no probability of an accident occurring and/or the severity of an accident if it does occur. Our framework supports Quiggin's (1992, 2002) contribution that comparative statics results are unchanged whether care efforts affect severity (**self-insurance**) or accident probability (**self-protection**), as defined in Ehrlich and Becker (1972), but we show that these features do affect the optimal design of subsidy policies, and the volume of transfers into the polluting industry that they require. Specifically, we show that self-protection forces the subsidy to pay in every period, even if no accident occurs, and we show that self-insurance forces the subsidy to pay for all reductions below a decoupled threshold, rather than paying for actual pollution reductions. In the agnostic case of either prevention mechanism being active, both prescriptions hold and the optimal subsidy amounts to an *ex ante* bribe for agents to participate in the strict liability regime.

In cases where high costs make subsidies infeasible, we propose a system of mandatory insurance which can still reach an optimal solution. This policy induces optimal choices of care and cleanup, but it keeps all monies within the polluting industry in expectation, so it is budget-neutral. The insurance policy thus has parallels to existing policies, like a carbon tax, whereby optimal behavior is induced but no payment is made to the parties suffering from global warming. In addition, we show that the within-industry distribution of resources under this policy is similar to that of tradable pollution permits; 'dirtier' firms end up subsidizing 'cleaner' firms according to the disparities between their optimal levels of pollution.

The next section introduces a general framework of decision-making under uncertainty (a principal-agent game), with the innovation that agents' actions before *and* after the accident can affect economic outcomes. Then, we develop optimal tax and subsidy policies, demonstrate the key results, and introduce a mutual insurance policy which is revenue-neutral for the regulator in expectation. We conclude with a discussion of policy implications and areas for future research.

We frame the accidents problem in terms of a principal-agent game between one regulator and a continuum of accident-prone agents, representing an industry. The regulator sets the rules of the game, recognizing the social cost of potential and realized accident damages, with the objective of maximizing total welfare, and the agents maximize expected profits. Each agent faces a known distribution of accident outcomes, F, an *a priori* choice of care to prevent accidents from occurring and/or decrease their severity, and a choice of cleanup efforts if an accident does happen. The agents are differentiated by type, which is a characteristic representing the inherent riskiness of their operations. For an oil company type might constitute onshore vs off-shore drilling, and for industry it may reflect production using more hazardous chemicals. As discussed above, the regulator cannot observe agents' actions (or type) or an accident's innate severity if it occurs, but both parties are aware of any damages, which are measurable and traceable to their source.

Our framework closely follows the notation of Hanley, Shogren and White (2007, p.401), herein "HSW". We start by considering the decisions of an individual agent, with riskiness type θ. Let w denote the benefits of economic activity, and let L denote a personal loss to the agent, which is increasing in accident severity, if an accident occurs. The *ex ante* choice of prevention expenditure is denoted z, which can be self-protection (lower probability of accident occurring), self-insurance (decreased severity if an accident does occur) or both. Prevention efforts might include avian flu vaccination by a poultry producer, the use of water sprayers and overflow/blow-off tanks at chemical plants, or measures to reduce the likelihood and severity of an oil spill.

If an accident does happen, the accident severity is parameterized by the continuous, non-negative random variable, $\gamma \in \Gamma = [0, \bar{\gamma}]$, which could be infection rate or viral count, or explosion or spill severity, depending on the application. The severity parameter could also indicate an accident happening in a sensitive location, as with a fire in dense apartment units, an oil spill in the habitat of an endangered species, an industrial accident in a city center, or a disease outbreak in a vulnerable population. We consider this parameter to be an increasing hazard, which is to say that a higher γ means a more harmful accident, and a realization of $\gamma = 0$ means no accident occurs. The accident severity follows a probability distribution, $\gamma \sim F(\gamma; \theta, z) : \Gamma \to [0,1]$, which is dependent both on an agent's type, and on his prevention efforts.

If an accident does occur, there is social damage $D(\gamma, a)$, from the spread of fire, disease or the release of hazardous substances, which is denominated in dollars. An agent can reduce the damage by an abatement expenditure, *a*, where $D_a < 0$, but with decreasing marginal returns, so $D_{aa} > 0$. Abatement measures can include assistance with fire control and cleanup of hazardous materials, slaughter of sick animals and sterilization of facilities and equipment, and notifying proper authorities to prevent the spread of harm. Clearly, the issue here is that cleanup efforts are costly to the accident victim, but they only benefit society. Accounting for social costs, as the agent will do when behaving optimally, yields the objective function:

$$E[\pi] = w - z - \int_0^{\bar{\gamma}} \left(L(\gamma) + D(\gamma, a(\gamma)) + a(\gamma) \right) dF(\gamma; \theta, z)$$

where E denotes the expectation operator and the integral is the Stieltjes integral. We assume that z and θ interact with F in the sense of **first-order stochastic dominance**, so that z makes the distribution unambiguously better and θ makes it unambiguously worse. Formally,

$$F(\gamma; \theta, z) \underset{FSD}{\succ} F(\gamma; \theta, z') : z < z' \text{ and}$$

$$F(\gamma; \theta, z) \underset{FSD}{\prec} F(\gamma; \theta', z) : \theta < \theta'$$

where dominance in this case implies a 'worse' distribution because our focus is the distribution of the social cost of accidents. Equivalently, we can say $F_z \geq 0 \ \forall \gamma$ and strictly greater for some γ. Similarly, $F_\theta \leq 0 \ \forall \gamma$ and strictly less for some γ. We also assume that prevention experiences decreasing marginal returns, so $F_{zz} \leq 0$, because damage-control spending uses up the lowest cost measures first.

While restrictive for ranking lotteries, the first-order stochastic dominance assumption is actually quite general when evaluating accident prevention mechanisms.

Recall the two common specifications of self-protection and self-insurance. Self-protection, where care affects only the probability of an accident, often models expected damages as $p(z)\cdot\int_{\Gamma} D\,dF(\gamma;\theta)$. Similarly, self-insurance, where care only affects severity when the accident does occur, might show damages modeled as $\int D(\gamma,z,a)\,dF(\gamma;\theta)$. Our agnostic model of generic first order stochastic dominance allows either or both of these specifications to hold.

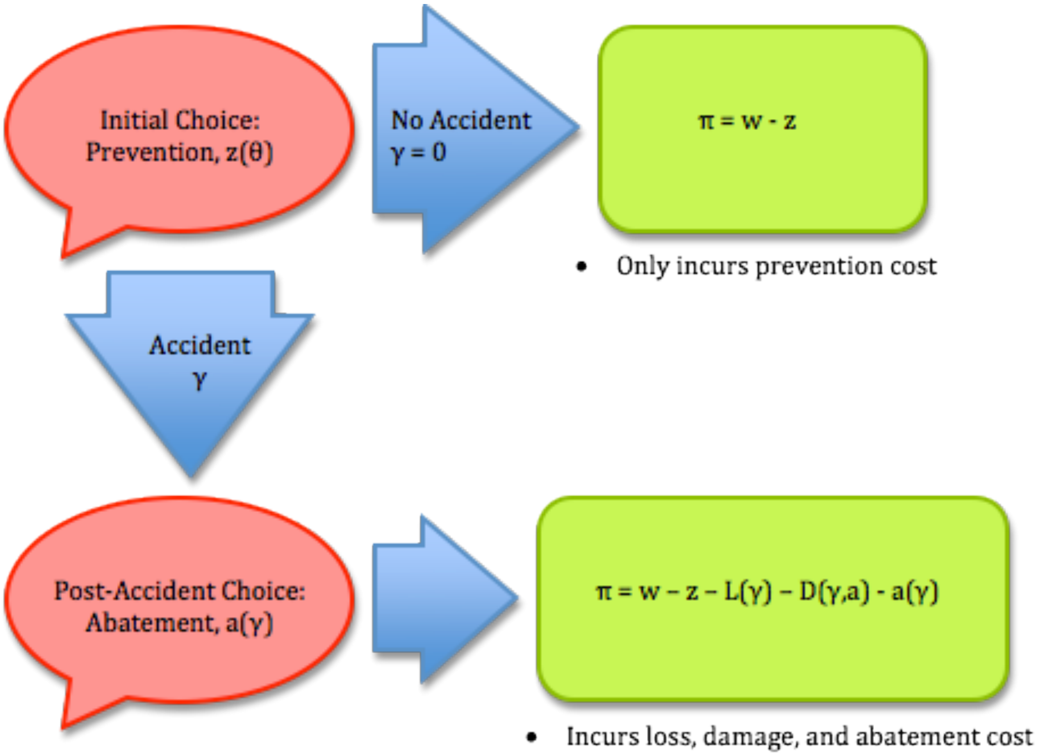

Figure 1 presents the risk-generation process captured by our framework. We solve for optimal behavior through a backwards inductive approach; first, we develop an optimal abatement response for every outcome if an accident does occur, and then we use this information to solve for the *ex ante* care choice. Since no accident occurs when $\gamma = 0$, the agent faces two basic scenarios:

$u = w - z$ with probability $F(0)$;

$u = w - z - L(\gamma) - D(\gamma, a) - a$ otherwise, where the accident scenario is dependent on the realized severity, γ . Thus, if an accident occurs, the agent solves:

$$\max_a \pi = w - z - L(\gamma) - D(\gamma, a) - a$$

where the optimal abatement expenditure is given by setting the marginal cost equal to the marginal benefits of damage reduction in the first order condition below:

$$-D_a = 1 .$$

The second order condition, $-D_{aa} < 0$, ensures that the optimal choice exists and is unique, as long as the initial benefits are greater than the marginal cost - which we assume is the case, so abatement activities are worthwhile at some level. Taking comparative statics results in the traditional fashion, we find that $a^*(\gamma)$ responds to changes in γ according to the sign of $-D_{\gamma a}$, about which we have made no assumption. For example, a more severe industrial fire might burn hotter, causing water sprayed on it to evaporate rapidly - an example of decreasing returns to the water expenditure. In essence, the cross-partial effect boils down to an empirical question - specific to the problem at hand - as to whether abatement experiences increasing or decreasing marginal returns when disasters are more severe.

Since the optimal containment response is deterministic, given the state of nature, we are equipped to evaluate the up front choice of care - with this information in mind. The *ex ante* choice is characterized by the following problem:

$$\max_z E[\pi] = w - z - \int_0^{\bar{\gamma}} \left(L + D^* + a^* \right) dF(\gamma;\theta,z)$$

where $D^* \equiv D\left(\gamma, a^*(\gamma)\right)$, and the other arguments are suppressed for clarity. The optimal care investment is given implicitly by the first order condition:

$$\frac{\partial E[u]}{\partial z} = -1 - \frac{\partial}{\partial z}\left[\int_0^{\bar{\gamma}} \left(L + D^* + a^* \right) dF(\gamma;\theta,z) \right] = 0$$

To verify an interior solution is possible, let $\tau(\gamma) \equiv L(\gamma) + D\left(\gamma, a^*(\gamma)\right) + a^*(\gamma)$ denote the social cost of an accident, which we know to be monotone increasing in γ by the envelope theorem. Then,

$$\begin{aligned} \frac{\partial}{\partial z}\int_0^{\bar{\gamma}} \tau(\gamma) dF(\gamma;\theta,z) &= \frac{\partial}{\partial z}\left[\tau(\gamma)\cdot F(\gamma;\theta,z)\Big|_0^{\bar{\gamma}} - \int_0^{\bar{\gamma}} F(\gamma;\theta,z) d\tau(\gamma) \right] \\ &= \frac{\partial}{\partial z}\left[\tau(\bar{\gamma}) - \int_0^{\bar{\gamma}} F(\gamma;\theta,z)\tau_\gamma(\gamma) d\gamma \right] \\ &= -\int_0^{\bar{\gamma}} F_z(\gamma;\theta,z)\tau_\gamma(\gamma) d\gamma < 0 \end{aligned}$$

where the first equality follows from integration by parts. The second equality follows from observing that $F(\bar{\gamma}) = 1 \ \forall z$ and converting the Stieltjes integral to a Riemann integral. Therefore, the first order condition amounts to setting:

$$\int_0^{\bar{\gamma}} F_z(\gamma;\theta,z)\tau_\gamma(\gamma)d\gamma = 1,$$

and accordingly, the second order condition for the care choice is:

$$\int_0^{\bar{\gamma}} F_{zz}(\gamma;\theta,z)\tau_\gamma(\gamma)d\gamma < 0.$$

As with the abatement choice, we assume that the marginal benefits of prevention are initially greater than the marginal costs, so care activities are worthwhile. Thus, the optimal prevention choice, *z*, which might be investment in avian flu vaccine, or emergency training for workers, is selected such that the marginal benefits of preventing harm to the producer and limiting externality exposure for others are set equal to the marginal cost of these efforts. The second order condition is everywhere negative, ensuring concavity, so the optimal choice z^* exists for each agent, and is unique.

Finally, there is a continuum of agents, of unit measure, who are differentiated by a type parameter, $\theta \in \Theta$, which is a continuous variable with probability distribution $G : \Theta \to [0,1]$. The type parameter can be considered as riskiness, or a propensity for more frequent and/or severe accidents. That is, a higher θ means a more risky agent. Riskiness might be a geographical element, like location within a city – as pertains to disease or fire risk, though we assume this parameter to be unobservable by the regulator, in general. Type could also index outdated equipment or the use of hazardous chemicals, or high interaction rates with other farms for the livestock producer.

For the care choice, we derive comparative statics results for the effect of increasing an agent's type, obtaining:

$$\frac{\partial z^*}{\partial \theta} = -\frac{\partial^2 E[u]}{\partial z \partial \theta} \Big/ \frac{\partial^2 E[u]}{\partial z^2} = \frac{\int_0^{\bar{\gamma}} F_{z\theta}(\gamma;\theta,z)\tau_\gamma(\gamma)\,d\gamma}{-\int_0^{\bar{\gamma}} F_{zz}(\gamma;\theta,z)\tau_\gamma(\gamma)\,d\gamma}.$$

So, optimal prevention will adjust to an agent's type according to the sign of $F_{z\theta}$. A traditional assumption is more riskiness will *increase* the marginal returns to care efforts, $F_{z\theta} > 0$, which is consistent with a Cobb-Douglas specification, with multiplicatively separable risk-generating functions from the environmental health literature, as in Starr (1985), and with exponential dose-response functions used in epidemiology (Wilson and Crouch, 1987; Bogen, 1995; Lichtenberg, 2010). However, this assumption can be controversial, as noted in HSW (2007, p.403), and is likely a problem-specific empirical question - as was the case in our discussion about damage containment efforts and accident severity, above. We leave for future research the examination of cases where type means more frequency but lower severity (i.e., not an FSD shift), as might occur when a metermaid has a high probability of car accidents in congested spaces, but collisions very often occur at low speed.

Given the above setup, we are able to characterize the social optimum as the maximum, expected aggregate welfare - with optimal choices by all agents, *ex ante*, and optimal response strategies *ex post*. Since π denotes individual profits, let $\prod$ denote the aggregate, so the maximal, expected total welfare is defined as:

$$E\left[\Pi\right]^{*}=\int_{\Theta}\left[w-z-\int_{0}^{\bar{\gamma}}\left(L+D^{*}+a^{*}\right)dF\left(\gamma;\theta,z^{*}\left(\theta\right)\right)\right]dG\left(\theta\right)$$

where the loss, damage and abatement terms are zero when no accident occurs, but are optimized if it does. This specification also allows for each agent to experience his own realization of γ , whose correlation across agents we have not yet addressed. Under our risk neutrality assumption, correlation would not affect the expectation operator, but future research may be needed to examine its effects when risk preferences are considered.

We now turn to developing policies which attain the social optimum; in essence, a new triumvirate of environmental policies under uncertainty. Before doing so, we turn briefly to the unregulated case for comparison.

The Unregulated Case

Without a liability standard, or other form of regulation, agents will choose sub-optimally *ex ante*, and there will be no *ex post* response to contain accidental harm to others.[7] This might occur if a livestock producer is not responsible for spreading swine flu, or if an oil refinery is not responsible for releasing air pollution. In the context of our framework, the unregulated case is suboptimal because the second stage disappears.[8]

Thus, the unregulated agents only solve the problem of making first-period prevention investments according to their own best interest:

$$\max_{z} E\left[\pi^{UR}\right] = w - z - \int_{\Gamma} L(\gamma)\, dF(\gamma;\theta,z)$$

$$\pi_z^{UR} = -1 + \int_0^{\bar{\gamma}} \left(F_z(\gamma;\theta,z) L_\gamma\right) d\gamma = 0 \quad .$$

As in the optimal case, the unique solution is defined implicitly by the first order condition. The agent *does* incur some personal loss as a result of the accident, so the prevention

---

[7] Neglecting, of course, any utility payoff from 'doing the right thing.'

[8] Consider the application of our framework to the poison-gas disaster at the Union Carbide pesticide plant in Bhopal, India in 1984, as documented by Eckerman (2005). The chemical manufacturer was operating essentially as if unregulated, as evidenced by the choices made. Care: safety measures were turned off, others left on were inadequate, and the operating crew was both under-trained and undermanned – all in order to save on the costs of these measures. Methyl isocyanate was also used instead of less hazardous, but more expensive alternatives. Containment: Once the accident occurred (a pressure spike, causing poison-gas to be released into the surrounding community), no efforts were made to notify authorities of the gas leak (it was denied at the moment of the disaster), and no assistance was provided to medical responders about the nature of the chemical exposure suffered by accident victims.

efforts still exist at a positive level in the absence of regulation, but no regulation means agents will under-prevent, relative to the social optimum.

Proposition U1: $z_{UR}^{*} < z^{*}$

Proof: Follows directly by inspection. First, fix the agent's type, θ. Then, by substitution of the first order conditions, we obtain:

$$\begin{aligned}
\int_0^{\bar{\gamma}} \left(F_z\left(z_{UR}^{*}\right) \cdot L_\gamma\right) d\gamma \qquad &= \int_0^{\bar{\gamma}} \left(F_z\left(z^{*}\right) \cdot \tau_\gamma\right) d\gamma \\
&= \int_0^{\bar{\gamma}} \left(F_z\left(z^{*}\right) \cdot \left(L_\gamma + D_\gamma^{*} + a_\gamma^{*}\right)\right) d\gamma \\
&> \int_0^{\bar{\gamma}} \left(F_z\left(z^{*}\right) \cdot L_\gamma\right) d\gamma
\end{aligned}$$

Since, $F_{zz} \leq 0 \ \forall \gamma$ and strictly less than zero for some γ, it follows that $z^{*} > z_{UR}^{*}$.

□

The preceding proposition verifies our earlier claim that some prevention will still exist - the polluting firm has its own losses to protect, after all - but that it will be lower than the optimal level because of a disregard for social accident costs. Unregulated behavior and the associated expected profits will also inform the participation decision when agents must be induced to accept the subsidy program.

Optimal Policy – Strict Liability/Penalty System

As with the point-source externality problems already assessed in deterministic frameworks, tax policy under uncertainty can maximize social welfare similarly by forcing polluters to "internalize the externality" – that is, they will account for the social cost of their actions as part of their decision-making process. In fact, we explicitly demonstrated this point in the derivation of the social optimum, above. The key element of stochastic pollution taxes is that the tax amount, or even the unit pollution tax, cannot be fixed in advance (as in Weitzman, 1974) because the optimal containment response is state-dependent - which would lead to second-best outcomes. Thus, we propose that a policy regime of *strict liability* will lead to socially optimal behavior by making agents liable to pay *D* after an accident, perfectly aligning their personal incentives with the social objective.

Some considerations are worth mentioning here. The liability system proposed here relies on perfect detection and traceability of the social damages, either by regulators or by the individuals affected, and no transaction costs of enforcement. However, it has been shown by Polinksy and Shavell (1992) that costs of detection and enforcement are really just a part of the externality, so these costs can be included in the damage function, *D*, without loss of generality. Even with the detection and traceability problem solved, other constraints might still exist; as van't Veld (1997) suggests, imposing high penalties may not be practical because of limited ability to pay - bankruptcies create an effective upper bound on financial penalties, preventing proper alignment of incentives for low probabilities of detection. This is the judgment-proof problem, which Innes (1999a) addresses in a liability setting by applying a stochastic penalty to balances the distribution of outcomes by mandating over-payment of fines above a certain threshold. This approach, of course, breaks down when optimal responses need to be maintained post-accident.

Thus, the problem of inability to pay can play a critical role in policy formation. While it might be argued that strict liability is the "purest" form of optimal regulation, since collected fines could theoretically be distributed to those harmed, many environmental taxes exist only to properly align incentives - there is often no mechanism for the actual payment of damages to those harmed. For example, carbon taxes in OEDC countries are not readily distributed to a farmer in Afghanistan who loses his farm due to global warming. Thus, for those unable to pay the accidental damages, alternative policies may be sought to induce optimal producer behavior.

Subsidy programs may appear when there is a historical right to pollute, when there is limited ability to pay, or when there is political power that interferes with enforcement of strict liability (Bulte et al, 2008). The term, *payment for environmental services,* reflects exactly this situation where individuals must be induced to voluntarily comply with environmental policy. We show in the next section that a compensation policy can still attain socially optimal outcomes, both *ex ante* and *ex post*, but that its design must fundamentally differ from its deterministic counterpart.

Compensation Policy – Abatement Subsidies

Consider a classical abatement subsidy in a deterministic setting. Depending on the amount of information available to the environmental regulator, agents can be compensated for actual abatement of damages or for curtailing production to limit emissions. In a stochastic environment, the optimal behavior is state dependent, so the *ex ante* care investment is important, but it is not sufficient for an optimal outcome. We demonstrate two special issues that arise: first, if care affects the probability of an accident then the subsidy must pay each agent in every period - even if no accident occurs. This requirement alone may test a regulator's budget constraint if accidents are rare.

Second, if care affects severity then the subsidy cannot pay compensation for actual abatement of damages - the payment must instead be based on abatement below a fixed (decoupled) threshold, e.g., $S = D_0 - D$. The decoupling requirement exists because compensation tied to actual abatement perverts incentives; it encourages risky behavior by rewarding agents when $D(\gamma, a = 0)$ is higher. There may also be practical measurement issues, since actual abatement is calculated based on un-contained damages, which may be unobservable, or may never come into existence.

Combining these two considerations, where care affects both probability and severity, or where the exact mechanisms are not known by the regulator, necessitates that agents be compensated for abatement below a fixed threshold, D_0, in every state of nature. Equivalently, agents may be paid an *ex ante* bribe to submit themselves to participation in a strict liability regime, preserving optimal incentives at the margin.

Note that there is nothing about this specification that requires the subsidy payment to be positive in all states of nature; for example, $D_0 = 0$ reduces to the optimal tax policy of the previous section. Therefore, the material distinction between the subsidy and tax policies is that a subsidy must induce voluntary participation from the accident-prone agents.

The first step in developing an optimal subsidy policy is recognizing that all agents must participate, because non-participation by any non-zero measure of agents means they make (suboptimal) unregulated choices. For the moment, we will assume universal participation to demonstrate the key results, and then evaluate the **participation constraint** explicitly in the context of the political-economic and budgetary implications of the subsidy policy under uncertainty.

For the following propositions, an agent's expected profits when participating in the subsidy program are given by:

$$
\begin{aligned}
E\left[\pi^{S}\right] &= w - z + \int_{0}^{\bar{\gamma}} \left(S - L - a^{*}\right) dF\left(\gamma;\theta,z\right) \\
&= w - z + \int_{0}^{\bar{\gamma}} \left(D_{0} - L - D^{*} - a^{*}\right) dF\left(\gamma;\theta,z\right)
\end{aligned}
$$

Proposition S1: Suppose D_0 is large enough to satisfy the participation constraint for all agents (universal participation), and suppose $S = D_0 - D : D \neq 0$, so that the subsidy <u>only</u> pays in states of nature where an accident occurs. If $dF(0)/dz > 0$, then the resulting choices will not be socially optimal for any agent.

Proof: This is a policy that pays for all abatement below a threshold, except when there is no accident. The expected profits for a subsidized agent are thus:

$$
\begin{aligned}
E\left[\pi^{S1}\right] &= w - z + \int_{\gamma \neq 0} \left(D_{0} - L - D^{*} - a^{*}\right) dF\left(\gamma;\theta,z\right) \\
&= w - z + D_{0} \cdot \left(1 - F\left(0;\theta,z\right)\right) - \int_{0}^{\bar{\gamma}} \left(L + D^{*} + a^{*}\right) dF\left(\gamma;\theta,z\right)
\end{aligned}
$$

where the variable arguments in the integrals are the same because losses, damages and abatement are zero when no accident happens ($\gamma = 0$). Maximizing expected profits of this form yields the first order condition:

$$\int_0^{\bar{\gamma}} F_z\left(\gamma;\theta,z_{S1}^*\right)\tau_\gamma\left(\gamma\right)d\gamma = 1 + D_0 \cdot F_z\left(0;\theta,z_{S1}^*\right)$$

where the right-hand side is less than one. Substituting in the first order condition for the social optimum, we see that this subsidy produces a suboptimal (lower) level of *ex ante* care because $F_{zz} < 0$:

$$\int_0^{\bar{\gamma}} F_z\left(\gamma;\theta,z^*\right)\tau_\gamma\left(\gamma\right)d\gamma = 1 < 1 + D_0 \cdot F_z\left(0;\theta,z_{S1}^*\right) = \int_0^{\bar{\gamma}} F_z\left(\gamma;\theta,z_{S1}^*\right)\tau_\gamma\left(\gamma\right)d\gamma$$

□

Proposition S2: Suppose D_0 is large enough to satisfy the participation constraint for all agents (universal participation), and suppose $S = D_0 - D$ in all states of nature if $dF(0)/dz > 0$. If $dF(\gamma)/dz > dF(0)/dz$ for some γ, then $D_0 = D(\gamma, a = 0)$ will not achieve the social optimum. That is, the subsidy threshold cannot pay for actual damage reductions due to abatement post-accident.

Proof: Expected utility is given by:

$$E\left[\pi^{S2}\right] = w - z + \int_0^{\bar{\gamma}} \left(D(\gamma,0) - L - D\left(\gamma, a^*(\gamma)\right) - a^*(\gamma)\right) dF(\gamma;\theta,z)$$

Taking a first order condition yields:

$$\int_0^{\bar{\gamma}} F_z\left(\gamma;\theta,z_{S2}^*\right)\left(\tau_\gamma(\gamma)-D_\gamma(\gamma,0)\right)d\gamma=1$$

$$\Rightarrow \int_0^{\bar{\gamma}} F_z\left(\gamma;\theta,z_{S2}^*\right)\tau_\gamma(\gamma)d\gamma=1+\int_0^{\bar{\gamma}} F_z\left(\gamma;\theta,z_{S2}^*\right)D_\gamma(\gamma,0)d\gamma>1=\int_0^{\bar{\gamma}} F_z\left(\gamma;\theta,z^*\right)\tau_\gamma(\gamma)d\gamma$$

which implies prevention will be smaller than optimal, exactly as in Proposition S1.

□

Combining the insights of Propositions S1 and S2 shows that when prevention activities affect both the probability and severity of accidents, or when the regulator does not know the exact, effective mechanism of prevention efforts, a subsidy payment will achieve socially optimal behavior if it pays in every period, based on abatement below a decoupled threshold. As discussed above, one equivalent method of decoupling payments in this fashion, and making sure they pay in every state of nature, is simply paying D_0 to each agent *ex ante* as compensation for participating in a system of strict liability for accidental environmental damages.

Now it remains to derive the conditions for fulfillment of the participation constraint, to ensure that all agents voluntarily participate in the subsidy program. Using the decoupled *ex ante* payment as a guide, Condition 1 says that the expected profits with no regulation must not exceed the expected profits generated by participation in the subsidy program, for all agents:

Condition 1: $E\left[\pi^S\right]\geq E\left[\pi^{UR}\right] \qquad \forall\theta.$

By comparing the objective functions, Condition 1 reduces to:

$$D_0 - \int D^* + a^* dF(\gamma;\theta,z^*) \geq z^* - z_{UR}^* + \int L\, dF(\gamma;\theta,z^*) - \int L\, dF(\gamma;\theta,z_{UR}^*)$$
$$\equiv \Delta z + \Delta E[L] > 0$$

where $\Delta z = z^* - z_{UR}^* > 0$, and the sign of the sum is known because z_{UR}^* minimizes the sum of $z + E[L \mid z]$, while $z^* > z_{UR}^*$ does not.

Thus, in expectation, the subsidy threshold must exceed the abatement expenditure and the social cost of accidents, $a^* + D^*$, which together we will dub the **total externality cost** ($TEC = D^* + a^*$). Put another way, Condition 1 shows that (in expectation) the subsidy payment, net of the expected abatement expenditure, must exceed the profits lost by an unregulated agent who voluntarily switches from the unregulated choice, z_{UR}^*, to the socially optimal z^*. In addition, we demonstrate a critical relationship of D_0 with the TEC in the following proposition. Namely, that the threshold will be bounded below by the expected, optimal-response TEC of the first-best care choice, z^*, and it will be bounded above by the hypothetical TEC corresponding to the unregulated prevention choice.

Proposition S3: If *the threshold* D_0 *is paid ex ante, and if the participation constraint (Condition 1) is satisfied with equality, then* $E\left[TEC^* \mid z^*\right] < D_0 < E\left[TEC^* \mid z_{UR}^*\right]$.

Proof: Let Condition 1 be satisfied with equality. Payment of D_0 *ex ante means it does not vary with the state of nature, so:*

$$
\begin{aligned}
& D_0 - \int D^* + a^* dF\left(\gamma;\theta,z^*\right) = \Delta z + \Delta E\left[L\right] \\
\Rightarrow \quad & D_0 = \Delta z + \Delta E\left[L\right] + E\left[TEC^* \mid z^*\right] > E\left[TEC^* \mid z^*\right] \\
& or \ = \Delta z + \Delta E\left[L\right] + \Delta E\left[TEC^*\right] + E\left[TEC^* \mid z^*_{UR}\right] < E\left[TEC^* \mid z^*_{UR}\right]
\end{aligned}
$$

where the Δ-notation follows that of Condition 1, and TEC^* denotes optimal-response with respect to abatement. The signs are known because the care choices minimize their respective loss-plus-cost functions.

□

Thus far we have established conditions for optimality of the subsidy, depending on universal adoption thereof, but it is important to note that this program has higher information requirements than strict liability. Specifically, the subsidy must pay from a threshold high enough such that the participation constraint is satisfied for all agents, so the regulator needs some information about how profits might change when switching from an unregulated to a regulated environment. However, any subsidy at the necessary level *or higher* will be sufficient to attain optimality - it will just do so at increased budgetary cost with a higher level of transfers into the polluting industry. Similarly, the subsidy program can thus solve the judgment-proof problem (when bankruptcies prevent the proper alignment of incentives) by raising D_0 to a level where polluting firms are always able to pay their liabilities.

To guarantee universal participation, it is clear that the regulator needs to know TEC information for the *highest cost* type, which may or may not be the riskiest (highest) type. Consider a threshold, $D_0(\theta)$, that satisfies the participation constraint with equality, by type. Applying the envelope theorem, we obtain:

$$\begin{aligned}\frac{\partial D_0(\theta)}{\partial \theta} \quad &= \frac{\partial}{\partial \theta}\left[\int_0^{\bar{\gamma}} F\left(\gamma;\theta,z_{UR}^*\right)\cdot L_\gamma \, d\gamma - \int_0^{\bar{\gamma}} F\left(\gamma;\theta,z^*\right)\cdot \tau_\gamma \, d\gamma\right] \\ &= \int_0^{\bar{\gamma}}\left[-L_\gamma \cdot \int_{z_{UR}^*}^{z^*} F_{z\theta}(\gamma;\theta,z)\,dz - \left(D_\gamma^* + a_\gamma^*\right)\cdot F_\theta\left(\gamma;\theta,z^*\right)\right]d\gamma\end{aligned}$$

where the first term has the opposite sign of $F_{z\theta}$ and the second term is positive because type leads to an objectively worse distribution of accident outcomes.

There are two effects, an **externality effect** and a **switching effect**, which arises from the impact of changing type and care on the expected personal loss. In general, $F_{z\theta} \leq 0$ is sufficient for the signs to agree and for the threshold to be increasing in type. However, we noted that the sign of this cross partial is an empirical question, and the traditional assumption (based on Cobb-Douglas specifications, etc.) is the opposite: $F_{z\theta} > 0$. Our intuition is that the externality effect dominates in many cases we care about, and though we do not demonstrate them rigorously, this intuition is supported by ideas like 'the externality and abatement cost are large relative to the personal loss,' and/or 'the marginal effect of type on the returns to care is small, relative to the effect of type on the externality distribution.'

So, if the participation threshold is everywhere increasing in type, then the regulator need only calculate the threshold for the highest one. If not, then he must know the 'worst' type, in terms of highest total externality cost and lost profits from regulation. Unfortunately, paying decoupled subsidies in every period, when subsidies must hold for all agents - and accordingly, for the most costly type - means outliers or skewness in the distribution of agents may make the optimal threshold very high relative to the participation constraint of most agents.

For example, chemical plants using highly toxic inputs may face higher externality costs in expectation, but the decoupled subsidy requires plants with less-toxic inputs to receive the same transfer, if the regulator cannot observe type. Thus, the information problems of the regulator may inflate the volume of subsidy payments to the polluting industry, or - to put it another way, the value of information about polluters is the direct reduction of information rents paid based on type. However, the participation constraint only binds on the expected *optimum* of total externality cost; the *optimal* total cost may be substantially lower than uncontrolled damages, as when prevention or abatement are inexpensive.

Thus far, we have considered two scenarios where the accidental externality problem is optimally regulated; either by a penalty system (or one of strict liability) that potentially places a heavy burden on accident victims, which might also comprise an industry with strong political influence, or by a compensation system where money flows into the polluting industry, and the volume of transfers increases for infrequent accidents (payment in every period), worse information, or more variation across producers' risk profiles.

In essence, a strict liability system is characterized by agents paying for what is known, but under a subsidy the regulator pays for what is *not* known, because the threshold has to cover the threshold for the worst type in order to obtain the optimal outcome. As a result, while the penalty policy may face ability to pay constraints from the polluters, the subsidy policy may face an ability to pay constraint for the government. Thus, we conclude our analysis by discussing the possibility of a mutual insurance policy, one that would be budget neutral for the regulator by keeping all funds within the polluting industry, in expectation.

Revenue Neutral Optimal Policy: Mutual Insurance

The third major tool of environmental policy is a system of cap and trade, or tradable pollution permits. While the state-dependent nature of stochastic externalities (when cleanup is possible), and realistic restrictions on government information about damage control technologies, prevent quantity-based environmental policy *per se*, we introduce a third policy to accomplish similar goals while attaining optimal behavior under uncertainty. The new policy involves mutual insurance for the polluters, while functioning like a system of bonded liability, and it retains the flavor of a system of tradable permits: revenue neutrality for the regulator (in expectation) because no funds move into, or out of, the polluting industry, and subsidization of 'safer' firms by riskier ones.

The insurance policy we propose functions similarly to the subsidy policy outlined in the previous section, but it is funded by the firms themselves and relies on the regulator imposing compulsory participation because there is no incentive from outside funds. Because participation cannot be made voluntary (which would require outside funding, as with the subsidy), the volume of transfers may be lower under the insurance program because abatement costs and forgone unregulated profits need not be compensated to achieve optimality.

As with the subsidy, the insurance program is structured to pay compensation below a specified threshold, D^I , in every period. However, since participation is mandatory, the threshold need not be tied to the highest type. Instead, our goal is assigning a threshold that leads to a balanced budget in expectation, where D^I and the premium, *P*, are the same for all agents. An agent's expected profits under the insurance program are thus given by:

$$E\left[\pi^I\right] = w - z - P + D^I - \int_0^{\bar{\gamma}} \left(L + D^* + a^*\right) dF \,.$$

The balanced budget constraint can be expressed as:

$$\int D^I - \int D^* \, dF\left(\gamma;\theta,z^*(\theta)\right) \, dG(\theta) = P$$
$$\Rightarrow D^I - P = \int\int D^* dF^* dG$$

which necessarily contains a degree of freedom between the threshold faced by the firms and the premiums collected - so either D^I or *P* can be established, and then the other calculated accordingly.

The key result of this program is that imposing a balanced budget means that all firms' net expected profits from participation (which may be negative) are based on their variation from the average expected level of environmental damage. To clarify this point, consider the insurance program where $P = 0$, so that:

$$D^I = \int\int D^* dF^* dG \,.$$

In this scenario, the expected profit function for each agent looks remarkably similar to that of the subsidy program outlined in the previous section. Each agent pays no premium up front and receives a decoupled subsidy payment of D^I in the *ex ante* decision-making stage. However, the level of the subsidy is set according to the *average* optimal environmental damage, so that some agents will exceed the threshold amount of damages and receive a negative subsidy (net loss), in expectation, while others will beat the threshold and experience a net gain. To complete the cycle, once accidents are realized *ex post*, all liability payments are paid into a common pool, leading to expected budget neutrality because some firms are net payers and others are net receivers of funds from the pool.

While there are distributional disparities among members of the industry, these differences are similar across all of our optimal stochastic externality policies, only subject to different baseline levels of wealth: the 'cleaner' or safer types are always better off, in expectation, than the riskier ones, when we define relative riskiness not by the type parameter necessarily, but by the optimal, expected TEC discussed in the subsidy section (which may be monotone in type, anyway). This process of subsidization from high-cost to low-cost agents mimics the results of deterministic cap and trade policies, where agents with higher costs of pollution control subsidize those who can carry out abatement more efficiently. However, unlike the outcome of tradable permits schemes this redistribution of resources is not an essential feature of the insurance policy design - instead, it is simply the by-product of the information problems we have assumed throughout.

If types are observable to the regulator, then each agent could be insured individually, leading to revenue neutrality for all, in expectation. Observability of types is the only such condition, though, because otherwise agents either collect or pay information rents when they participate in a compensation policy. To see this, consider the expected profits of an agent, θ, operating under the mutual insurance program, in the special case where $P = 0$ and $D^I = D^I(P = 0)$:

$$\begin{aligned}
E\left[\pi^{I*(\theta)}\right] &= w - z^*(\theta) + D^I - \int_0^{\bar{\gamma}} \left(L + D^* + a^*\right) dF\left(\gamma;\theta,z^*(\theta)\right) \\
&= w - z^*(\theta) + \int \left[\int_0^{\bar{\gamma}} D^*\, dF\left(\gamma;\theta,z^*(\theta)\right)\right] dG(\theta) - \int_0^{\bar{\gamma}} \left(L + D^* + a^*\right) dF\left(\gamma;\theta,z^*(\theta)\right)
\end{aligned}$$

Thus, outside the private costs experienced by each agent, the net cost (or profit) from participation in the subsidy program is given by the information rents - the departure of expected damages, $E\left[D^*\right]$, from the industry average. Each agent's net expected disbursement *from* the insurance pool is given by his **information rent**, resulting from the regulator's inability to observe his type, which is given by:

$$\begin{aligned}
R(\theta_0) &= \int \left[\int_0^{\bar{\gamma}} D^*\, dF\left(\gamma;\theta,z^*(\theta)\right)\right] dG(\theta) - \int_0^{\bar{\gamma}} \left(D^*\right) dF\left(\gamma;\theta_0,z^*(\theta_0)\right) \\
&= \int \left[\int_0^{\bar{\gamma}} D^*\, dF\left(\gamma;\theta,z^*(\theta)\right) - \int_0^{\bar{\gamma}} \left(D^*\right) dF\left(\gamma;\theta_0,z^*(\theta_0)\right)\right] dG(\theta) \\
&= \int \left[\int_\theta^{\theta_0} \left(\int_0^{\bar{\gamma}} D_\gamma^* \cdot F_\theta\left(\gamma;t,z^*(t)\right) d\gamma\right) dt\right] dG(\theta)
\end{aligned}$$

where the final equality results from integration by parts and simplifying, as above. Next, we implicity define the "average" agent in the sense of expected *optimal* environmental damages, so $\hat{\theta} = \theta : R\left(\hat{\theta}\right) = 0$. Thus, “above-average” agents are riskier and will be net payers into the pool, in expectation, while safer firms of below average riskiness will experience a net expected profit:

$$\theta \geq (<) \hat{\theta} \Leftrightarrow R(\theta) \geq (<) 0 ,$$

because $F_\theta < 0$ and the integral runs in the negative direction for $\theta > \hat{\theta}$. The resulting resource allocation is analogous to that of cap-and-trade policies for deterministic externalities - dirtier firms subsidize cleaner ones, defined relative to the industry average, because they do more environmental damage - even when their behavior optimally accounts for the social costs of production.

Discussion

The value of information about polluters is apparent - hidden types translate directly to information rents when policy is constrained to a compensation mechanism like subsidies or insurance. As a result, the budgetary cost, or the necessary volume of total transfers, of compensation-based policies can be lowered when polluters’ inherent riskiness is known. Another benefit of exposing types is allowing for adequate handling of entry into the polluting industry, which we have not addressed.

In fact, similar to the deterministic abatement subsidy, our proposed subsidy and insurance programs rely on the industry excluding firms when their optimal behavior is non-operation. When types are unobservable and/or when we consider the long run, where entry is possible, then these policies fall short just like their deterministic counterpart because firms will enter who operate at a net social loss. Thus, while robust to moral hazard issues, these policies are not robust to adverse selection, because these firms should not be producing at all - but they join the industry because the insurance or subsidy program allows for positive expected profits. Unfortunately, without being able to identify types, or impose barriers to entry, only the harshest policy of strict liability will ensure the proper composition of a stochastically polluting industry.

We foresee other areas where expanding the analysis may be helpful. While Just and Zilberman demonstrate the asymmetry of taxes and subsidies due to risk aversion, deriving the optimal insurance policy may be of interest in this setting, or when considering polluting firms who are loss averse. Risk preferences will also play a role in evaluating the accident mechanism itself. Correlation across agents - such as might happen over space, when accidents are weather-related - doesn't affect the expectation under risk neutrality, but it might affect the agents' welfare if they are utility maximizers, especially if their accidents harm one another. We leave this exploration as an area for future research.

Under risk neutrality, we have outlined three efficient policy regimes to deal with uncertainty, which fulfill the roles of their deterministic counterparts - especially in the broader sense of resource transfer into, out of, or within the polluting industry. Only the strict liability regime can produce funds to compensate outside accident victims, because it takes money from the industry as fines collected by the regulator - or even in the form of direct claims by those harmed. However, in many cases, no such mechanism readily exists, and this policy may be constrained by equity considerations or ability-to-pay constraints because the polluter is a victim who also adheres to the optimal standard of care.

The compensation policy can circumvent agents' ability-to-pay constraints, but may induce a constraint on the regulator's ability to pay, due to the volume of payments when accidents are infrequent, or when there are known high-risk outliers who cannot be identified (because they increase the decoupled threshold, which forms the basis of payment to all firms). While strict liability forces agents to pay for what is known (the damages), the stochastic abatement subsidy makes government pay for what is not known - types, actions, control technologies, etc. - resulting in large transfers into the polluting industry, and no mechanism for compensation of outside victims.

Like the other policies, the mutual insurance policy we propose preserves the marginal incentives of strict liability, but it recycles the fines into a pool so that all funds remain in the industry. While this mechanism denies compensation to the victims of environmental accidents outside the polluting firms, it is budget-neutral in expectation, and it facilitates a transfer from riskier to safer firms, as can occur under cap-and-trade in the deterministic setting.

Major environmental accidents do not always occur in a well regulated environment, but we have shown that classical, deterministic environmental economic policies can be adapted to uncertainty in recognizable forms. Their stochastic counterparts have similar distributional implications with respect to the polluting industry, though uncertainty exacerbates the disparities between resource allocations generated by the various policies. Our proposed risk-pooling scheme thus reflects a new policy ideal; namely, goals of efficiency and minimal redistribution of resources can be achieved simultaneously, even in the context of asymmetric or limited information and stochastic mishap.

References

William J. Baumol and Wallace E. Oates. The Theory of Environmental Policy. Cambridge University Press, New York, 1988.

Beard, T. R. (1990). Bankruptcy and care choice. *Rand Journal of Economics*, 21(4), 626–634

Kenneth T. Bogen. Methods to approximate joint uncertainty and variability in risk. *Risk Analysis*, 15(3):411–419, 1995.

Buchanan, James M. and Gordon Tullock. “Polluters' Profits and Political Response: Direct Controls versus Taxes.” The American Economic Review, Vol. 65, No. 1 (Mar., 1975), pp. 139-147

Bulte, Erwin H., Leslie Lipper, Randy Stringer and David Zilberman. "Payments for ecosystem services and poverty reduction: concepts, issues, and empirical perspectives." *Environment and Development Economics* (2008), 13:245-254

R.H. Coase. The problem of social cost. *Journal of Law and Economics*. 3 (1960) 1–44.

Cooter, Robert, and Ariel Porat. Does Risk to Oneself Increase the Care Owed to Others? Law and Economics in Conflict. *The Journal of Legal Studies*, 2000 29:1, 19-34

Ingrid Eckerman. The Bhopal Saga: causes and consequences of the world's largest industrial disaster. Universities Press (2005), 283 pages. ISBN: 8173715157, 9788173715150.

Aaron S. Edlin, Efficient standards of due care: Should courts find more parties negligent under comparative negligence?, *International Review of Law and Economics*, Volume 14, Issue 1, March 1994, Pages 21-34

Ehrlich, I. and G.S. Becker. “Market Insurance, Self-Insurance, and Self-Protection.” The Journal of Political Economy, Vol. 80, No. 4 (Jul. - Aug., 1972), pp. 623-648.

S.F. Hamilton and D. Zilberman. Green markets, eco-certification, and equilibrium fraud. *Journal of Environmental Economics and Management*, 52(3):627–

644, 2006.

Hanley, N., Shogren, J. and B. White (2007). Environmental Economics in Theory and Practice, Second Edition. Palgrave Macmillan: New York, NY.

Holmstrom, B. “Moral Hazard in Teams.” *The Bell Journal of Economics*, Vol. 13, No. 2 (Autumn, 1982), pp. 324-340

Innes, R. (1999a). Optimal liability with stochastic harms, judgement-proof injurers, and asymmetric information. *International Review of Law and Economics*, 19, 181–203.

Innes, R. (1999b). Self-policing and optimal law enforcement when violator remediation is valuable. *Journal of Political Economy*, 107, 1305–1325.

Just, R., & Zilberman, D. (1979). Asymmetry of taxes and subsidies in regulating stochastic mishap. *The Quarterly Journal of Economics*, 93, 139–148.

William M. Landes and Richard A. Posner. Causation in Tort Law: An Economic Approach. *The Journal of Legal Studies*, Vol. 12, No. 1 (Jan., 1983), pp. 109-134

Erik Lichtenberg and David Zilberman. Efficient regulation of environmental health risks. *The Quarterly Journal of Economics*, 103(1):167–178, 1988.

Erik Lichtenberg. “Economics of Health Risk Assessment.” *Annual Review of Resource Economics*, 2010. (*forthcoming*)

Katrin Millock, David Sunding and David Zilberman. Regulating Pollution with Endogenous Monitoring. *Journal of Environmental Economics and Management*, 44: pp221 - 241 (2002)

Mattias K. Polborn. Mandatory insurance and the judgment-proof problem. *International Review of Law and Economics*, Volume 18, Issue 2, June 1998, Pages 141-146

A. Mitchell Polinsky. Strict Liability vs. Negligence in a Market Setting. *The American Economic Review*, Vol. 70, No. 2, Papers and Proceedings of the Ninety-Second Annual Meeting of the American Economic Association (May, 1980), pp. 363-367.

A. Mitchell Polinsky and Steven Shavell. Enforcement Costs and the Optimal Magnitude and Probability of Fines. *Journal of Law and Economics*, Vol. 35, No. 1 (Apr., 1992), pp. 133-148

A. Mitchell Polinsky and Steven Shavell. Should Liability Be Based on the Harm to the Victim or the Gain to the Injurer? *Journal of Law, Economics, & Organization*, Vol. 10, No. 2 (Oct., 1994), pp. 427-437

Quiggin, J. "Risk, Self-Protection, and Ex Ante Economic Value: Some Positive Results." *Journal of Environmental Economics and Management* (1992) 23: pp.40-53.

Quiggin, J. "Risk and Self-Protection: A State-Contingent View." Journal of Risk and Uncertainty (2002) 25:2 pp.133-145.

Ribaudo, Marc, and Margriet Caswell. "Environmental Regulation in Agriculture and the Adoption of Environmental Technology." In *Flexible Incentives for the Adoption of Environmental Technologies in Agriculture*. Frank Casey, Andrew Schmitz, Scott Swinton, and David Zilberman, eds., pp. 7-25. Dordrecht: Kluwer Academic Publishers, 1999.

Kathleen Segerson. Uncertainty and incentives for nonpoint pollution control. *Journal of Environmental Economics and Management*, 15(1):87–98, 1988.

Shavell, Steven. Strict Liability versus Negligence. *The Journal of Legal Studies*, Vol. 9, No. 1 (Jan., 1980), pp. 1-25

Steven Shavell. Uncertainty over Causation and the Determination of Civil Liability. *Journal of Law and Economics*, Vol. 28, No. 3 (Oct., 1985), pp. 587-609

Shavell, Steven. The judgment proof problem. *International Review of Law and Economics*, Volume 6, Issue 1, June 1986, Pages 45-58

Shavell, S. (1987). Economic analysis of accident law. Cambridge: Harvard University Press.

Starr, Chauncey. 1985. "Risk Management, Assessment, and Acceptability." *Risk Analysis* 5(2): pp.97-102.

Swierzbinski, J.E., 2002. Guilty until Proven Innocent –Regulation with Costly and Limited Enforcement. *Journal of Environmental Economics and Management* 27(2), 127-146.

Klaas T. van't Veld. The judgment proof opportunity. CUDARE Working Paper No. 823, University of California, Berkeley, 1997.

Weitzman, Martin L. "Prices vs. Quantities." *The Review of Economic Studies*, Vol. 41, No. 4 (Oct., 1974), pp. 477-491

R. Wilson and E. A. Crouch. Risk assessment and comparisons: An introduction. *Science*, 236(4799):267–270, 1987.

Joshua Graff Zivin and Arthur A. Small. Risk sharing in Coasean contracts. *Journal of Environmental Economics and Management*, 45 (2003) 394–415

Joshua Graff Zivin, Richard E. Just, David Zilberman. Risk Aversion, Liability Rules, and Safety. *International Review of Law and Economics*, Volume 25, Issue 4, December 2005, Pages 604-623