ESSAYS

A NEW MODEL OF ADMINISTRATIVE ENFORCEMENT

Robert J. Jackson, Jr.[*] *and David Rosenberg*[**]

THIS Essay proposes a new method of monitoring regulatory compliance by a firm that operates multiple sources of risk, such as air polluting smokestacks. The expense of individually monitoring such sources may consume a large share of the agency's enforcement budget, undermining deterrence objectives. Under our approach, regulators would instead randomly select one of the firm's sources of risk, determine the firm's liability at that source, and apply that outcome perforce as determinative of liability at all of the sources. This method, which we call single-outcome sampling ("SOS"), replicates or improves deterrence generated by the current source-by-source enforcement model, but at a fraction of the cost. To demonstrate these benefits, we apply SOS to the EPA's monitoring of compliance with Clean Air Act regulations. We also address potential risk-bearing and judgment-proof costs associated with our proposal and explain how both problems can be solved.

INTRODUCTION

This Essay develops a novel method of enforcing administrative regulations. Our proposal will allow agencies to reduce the cost of monitoring regulatory compliance by a firm that operates multiple sources of risk, such as air-polluting smokestacks. The expense of individually monitoring such sources—a common strategy we refer to as

[*] Terence M. Considine Research Fellow in Law and Economics, Harvard Law School.

[**] Lee S. Kreindler Professor of Law, Harvard Law School. For helpful discussions and comments, the authors thank Michael Abramowicz, Sarah Abramowicz, Andrew S. Buzin, Dennis Connolly, Louis Kaplow, Aaron M. Katz, Ariel Porat, Christina J. Saunders, Peter Schuck, Steven Shavell, Richard C. Squire, James P. Sullivan, Jonathan B. Tager, Zachary D. Tripp, and participants in the Law and Economics Seminar at Harvard Law School. Jackson also gratefully acknowledges the Olin Center for Law, Economics, and Business for generous financial support.

"iterative monitoring"—may consume a large share of an agency's enforcement budget. At some point, budgetary constraints may force the agency to choose between allocating resources to iterative monitoring or to other enforcement priorities. If both courses of action would yield a net enforcement benefit, the agency's choice of one over the other necessarily sacrifices socially desirable deterrence.[1]

To remedy this problem, we propose a new method of random sampling that an agency can use in place of iterative monitoring to achieve the same level of deterrence at a fraction of the cost.[2] Specifically, under our proposal regulators would randomly select for monitoring only *one* source from among the total number that the firm operates. Upon determining the amount of liability at the selected source, the agency would apply that outcome perforce as determinative of liability for all regulated sources—the unselected sources as well as the selected source.[3] Our method confronts the firm with the same aggregate expected liability as it would face if all

[1] For discussion of how and why enforcement costs compromise deterrence objectives, see A. Mitchell Polinsky & Steven Shavell, Enforcement Costs and the Optimal Magnitude and Probability of Fines, 35 J.L. & Econ. 133, 135–39 (1992). See also Colin S. Diver, The Assessment and Mitigation of Civil Money Penalties by Federal Administrative Agencies, 79 Colum. L. Rev. 1435, 1466–68 (1979). In a recent example of the effect of reductions in enforcement budgets on deterrence, the FDA has requested a $10 million increase in funding in response to assertions that fewer inspections of food processors contributed to an outbreak in food-borne illnesses. See FDA, Fiscal Year 2008 Strengthening Food Safety Budget Request, available at http://www.fda.gov/oc/oms/ofm/budget/2008/BIB/PDF/5-FoodSafetyBCP(POM).pdf. And, even before manufacturers recalled Chinese imports due to product-safety concerns, a Commissioner of the Consumer Product Safety Commission testified before the Senate that reductions in product-safety enforcement staff had resulted in an "inability to have constant hands-on supervision [that] can result in products entering this country that do not meet U.S. safety standards." Thomas H. Moore, Comm'r, U.S. Consumer Prod. Safety Comm'n, Statement Submitted to Senate Committee on Commerce, Science, and Transportation (Mar. 21, 2007), available at http://www.cpsc.gov/pr/moore2007.pdf.

[2] The principal focus of the Essay is on regulatory deterrence by imposing financial liability on the firm for violating a negligence-type cost-benefit standard or for engaging in some activity subject to strict liability. For our purposes, "liability" encompasses taxes, fines, compensatory damages, and other forms of monetary assessments. We do not consider other modes of law enforcement such as imprisonment, seizures, and injunctions. "Compliance" refers to any conduct or response by the firm that the agency seeks to encourage through threatened liability.

[3] Our proposal would apply the determination of liability at the selected source only to sources controlled by the firm that the agency would otherwise subject to iterative monitoring. For convenience, we refer to these as "regulated sources."

of its regulated sources were monitored iteratively. Both strategies thus give the firm the same incentives for compliance. But, in requiring the agency to determine liability at only one source, our proposal achieves this deterrence result at a fraction of the cost of iterative monitoring.

Consider a firm operating a fleet of five trucks subject to a regulation requiring that all trucks have certain registrations. Assume that an agency would impose liability of $100 upon the firm for each unregistered truck and that two of the firm's trucks are unregistered. The agency could discover and sanction these violations by iteratively monitoring the firm's entire fleet: inspecting each truck for the appropriate registration and, upon finding two unregistered vehicles, levying total liability of $200. Instead, under our proposal, the agency could randomly select a single truck, assess the firm's liability for that truck, and apply that outcome as determinative of liability for all five trucks. If the agency selected an unregistered truck, the firm would bear total liability of $500 ($100 times five trucks); if the agency selected a registered truck, the firm would bear total liability of $0. Notably, using our approach, the firm would face the same aggregate expected liability of $200—a 40% chance of $500 in total liability and a 60% chance of $0 in total liability—as it would under iterative monitoring, but the agency would only have to inspect one truck to produce that result.

Our proposal, which we call "single-outcome sampling" ("SOS"), can replace iterative monitoring ("IM") without compromising deterrence. The key to this result is that SOS confronts the firm with the same aggregate expected liability as it would face under iterative monitoring. The basic case for our proposal is that it produces the same level of deterrence as iterative monitoring, but at a much lower cost.[4]

SOS is readily applicable to a variety of regulatory regimes. Agencies can use SOS whether the criterion for determining liability is strict liability or a negligence-type cost-benefit analysis. SOS also applies whether iterative monitoring would induce the firm to standardize or customize compliance measures. The advantages of

[4] Our assessment of SOS realistically assumes that reducing the number of monitored sources will reduce monitoring costs. SOS may well enable an agency to redeploy the surplus resources to enhance enforcement. See infra Part II.

SOS hold whether the regulatory requirement is set for each source individually or for all sources in the aggregate. Further, in addition to reducing the number of sources of risk monitored at a particular point in time, SOS can also be used to reduce the frequency of monitoring.

The single-source model of sampling is theoretically sound as well as practically workable. SOS employs ordinary processes of random selection by which each member of a population is given an equal chance of being chosen. It comports with the standard theory of optimal law enforcement, according to which it may be socially optimal to reduce enforcement cost by lowering the probability of detection while raising the magnitude of sanctions to maintain deterrence.[5]

Because of its straightforward mechanism, SOS promises greater efficiencies than the strategies agencies currently use to reduce the cost of iterative monitoring. Regulators principally rely on statistical sampling or offsetting reduced enforcement budgets with higher sanctions, but these methods are more structurally complex than SOS and consequently are likely to be more costly. In contrast to SOS, statistical sampling requires the agency to derive an average result as the basis for extrapolating liability to all sources.[6] To assure a statistically reliable basis for establishing aggregate liability, the agency must

[5] See Gary S. Becker, Crime and Punishment: An Economic Approach, 76 J. Pol. Econ. 169, 170 (1968). For elaboration of the theory and summary of the leading developmental and critical literature, see Steven Shavell, Foundations of Economic Analysis of Law 479–84, 491 (2004), and A. Mitchell Polinsky & Steven Shavell, Public Enforcement of Law, *in* 3 The New Palgrave Dictionary of Economics and the Law 178 (Peter Newman ed., 1998). An early application of the theory in the administrative context was provided in Diver, supra note 1, at 1467. For an application of the approach in the civil liability context, see A. Mitchell Polinsky & Yeon-Koo Che, Decoupling Liability: Optimal Incentives for Care and Litigation, 22 RAND J. Econ. 562, 562 (1991).

[6] For example, federal agencies enforce laws against racial discrimination in employment by barring firms from receiving government contracts when a sampling of representative employment decisions yields a statistically significant variance between the racial composition of the applicant pool and the actual composition of the firm's workforce. 28 C.F.R. § 50.14(4)(D) (2007); Office of Federal Contract Compliance Programs, U.S. Dep't of Labor, Federal Contract Compliance Manual 239–41 (1993). The Department of Agriculture imposes liability for violations of product-quality regulations based on a statistical average. See, e.g., 68 Fed. Reg. 1556, 1562 (Jan. 13, 2003) (to be codified at 7 C.F.R. pt. 723) (assessing liability for prohibited tobacco sales based on the average market price in the preceding year). Courts have also approved the use of statistical sampling to determine average liability to govern a larger population of claims. See, e.g., Hilao v. Estate of Marcos, 103 F.3d 767, 786–87 (9th Cir. 1996) (approving the use of a random sample of class claimants to establish damages for the class).

evaluate the sources involved and determine the appropriate size and composition of the sample. SOS, in contrast, establishes aggregate liability without the complication and cost of statistical sampling. Similarly, budgetary constraints frequently compel agencies to limit or reduce enforcement expenditures and offset lost deterrence by raising sanctions. It is far from straightforward, however, for the agency to determine how much the decrease in enforcement expenditures will lower the probability of detection as well as how much the increase in the level of liability will raise deterrence.[7] Under SOS, by comparison, the ratio between the probability of detection and the magnitude of sanction adjusts automatically—and precisely.[8] SOS thus offers agencies the means to duplicate existing incentives for compliance while avoiding the informational and decisional burdens of alternative monitoring strategies.[9]

[7] See Shavell, supra note 5, at 484 (noting that the significant difficulty in estimating the deterrence shortfall may require regulators to make rough judgments).

[8] Agencies could, of course, implement the lower-detection, higher-sanction policy with more mathematical precision. Regulators could, for example, randomly inspect each regulated source and, to offset lost deterrence, multiply liability by the inverse of the probability that the source will not be inspected. If the probability that each of a firm's two sources will be inspected is 50%, then the probability multiplier would be two. In this example, the strategy is indistinguishable from SOS. This is true, however, only because the probability multiplier happens to equal the number of regulated sources. If, as observation suggests is likely, the agency fixes the probability of inspection of a given source without regard to the number of regulated sources, then—even though the probability multiplier appropriately offsets lost deterrence—the strategy will be less cost effective than SOS. When the probability multiplier is *larger* than the number of sources, the strategy will impose greater risk-bearing and judgment-proof costs than would SOS; when the multiplier is *smaller* than the number of sources, the strategy will commit the agency to expending more enforcement resources than would SOS, with the excess increasing as the number of sources rises. In further contrast, SOS has the added efficiency of allowing the agency to impose liability without first determining, possibly at substantial cost, the total number of regulated sources. When the agency finds no liability at the selected source, it never incurs that burden. Even when the agency does hold the firm liable at the selected source, the agency may avoid any extra expenditure to identify and tally the remaining regulated sources if the firm is otherwise motivated to disclose them in due course, for example, by tax or licensure laws, such as when a trucking company must pay outstanding parking fines to re-register its vehicles. This advantage of SOS—that agencies can use it without determining beforehand or possibly at any time the firm's total number of sources—is the reason we describe our proposal as requiring regulators to *apply* the determination of liability at one source to all sources rather than to *multiply* liability by the total number of sources.

[9] We proceed on the assumption that SOS provides a competitive alternative to imposing average liability derived from statistical sampling or inflating the sanction to offset reduced enforcement budgets. As we have explained, there are several reasons

In Part I, we describe the mechanics of SOS and its various basic applications. To demonstrate the potential benefits of SOS, in Part II we apply the proposal in the context of the EPA's monitoring of compliance with Clean Air Act regulations. The EPA relies on direct inspection of stationary sources of air pollution to determine whether firms are in compliance with emissions standards, but the agency monitors far fewer sources with far less frequency than Congress authorized and EPA regulations mandate.[10] We show that, by adopting SOS, the EPA could not only secure the same compliance at lower cost, but also substantially expand its regulatory program using existing resources.

In Part III, we consider potential costs of applying SOS to cases in which SOS might confront a firm with the possibility of incurring substantial total liability. In particular, this prospect might increase the firm's risk-bearing cost (to the extent that the firm is risk averse) or decrease its incentives for compliance (to the extent that the firm anticipates that it will be judgment-proof).[11] We explain that the agency can address these problems by permitting wider sampling of the firm's regulated sources. We conclude by briefly describing extensions of our proposal to a variety of regulatory regimes and to the civil liability context.

I. THE SOS MECHANISM

We begin by explaining how SOS replicates the level of deterrence achieved by IM enforcement of strict liability regulations. We then demonstrate that the benefits of SOS extend to the negligence context.[12] Notably, we show that the symmetry in deterrence

to think that SOS may be the lower-cost alternative in many administrative contexts. Whether this would be true in any particular administrative context is an empirical question that we do not attempt to resolve here.

[10] See infra Part II and sources cited therein.

[11] For a discussion of the potential for enforcement policies that, like SOS, lower the probability and magnify the degree of sanction to pose risk-bearing and judgment-proof problems, see Polinsky & Shavell, supra note 5 (summarizing the literature on this point).

[12] Negligence and strict liability refer to the basic rules that agencies and courts use to determine the extent to which the target of regulation should bear the social costs of its risktaking conduct. For purposes of our discussion, the negligence rule conditions liability on a finding that the risktaking conduct violates a prescribed standard of care (or ratio of social cost to benefit). In its basic formulation, strict liability ignores the nature of the conduct at issue, imposing liability on the fact of risk or harm alone. For simplicity, in this discussion we do not distinguish between a negligence rule that

between IM and SOS holds whether negligence regulations call for standardized or customized precautions. We also show that SOS can function whether the agency sets a negligence standard of compliance for each source individually or for all sources on an aggregate basis. Finally, we explain how SOS can be used to further lower enforcement costs by reducing the frequency with which sources are monitored within a given period of time.[13]

In general terms, under SOS the probability that the agency will monitor a regulated source is equal to $1/n$, where n represents the number of sources otherwise subject to monitoring under IM. If a firm operates two regulated sources, the probability that each will be monitored under SOS is 50%. But the total assessment for noncompliance is multiplied by n to keep expected liability constant. The key is that, despite reduced monitoring, the firm's incentives for compliance remain the same because liability increases or decreases by a constant factor determined by the number of regulated sources controlled by the firm.

A. Strict Liability

Strict liability provides a straightforward basis for demonstrating how SOS replicates the aggregate expected liability produced by IM. While offering notable advantages over liability based on a negligence-type cost-benefit analysis,[14] strict liability may be underused because of its relatively high enforcement costs.[15] By reducing those costs, SOS may render the strict liability approach more attractive to regulators.[16]

specifies a standard of care *before* (a "command-and-control" rule) or *after* (an "ex post" rule) the risktaking conduct occurs; nor do we distinguish among negligence or strict liability regulations that assess liability based on risk versus actualized harm. None of these variations affects our analysis or conclusions.

[13] For additional discussion and an important caveat, see infra Section I.C.

[14] See, e.g., A. Mitchell Polinsky, An Introduction to Law and Economics 50–52 (3d ed. 2003). The strict liability approach, for example, entails lower information costs, obviating the need for determining standards to appraise the reasonableness of risky conduct. Shavell, supra note 5, at 229. Strict liability is also more cost effective in securing socially appropriate adjustments to the level of risky activity. Id. at 197–99.

[15] Shavell, supra note 5, at 283. We note, however, that by design (for example, an objective standard of care) or mistake (for example, an excessive, impractical to satisfy standard of care), the negligence rule often operates in some degree as a rule of strict liability.

[16] For a discussion of this point in the context of environmental regulation, see infra notes 50–51.

To illustrate how SOS produces the same aggregate expected liability as IM under strict liability, consider a firm operating two factories, *A* and *B*, producing 10 and 20 units of a regulated chemical, respectively. Assume that, under IM, the agency monitors the output at each factory and holds the firm strictly liable in the amount of $1 for each unit of the chemical produced. The following table describes the aggregate expected liability imposed on the firm under IM:

Facility	*Units Produced*	*Expected Liability*
A	10	$10
B	20	$20
Total	**30**	**$30**

SOS would replicate this aggregate expected liability. Under SOS, if source *A* is selected, the firm expects aggregate liability of $20; if source *B* is selected, the firm expects aggregate liability of $40.[17] Expecting each result with 50% probability, the firm anticipates aggregate liability of $30—the same as under IM. Facing the same aggregate expected liability under IM and SOS, the firm has the same incentives to comply under both regimes.

B. Negligence

Negligence is a ubiquitous administrative standard for determining liability and generally entails a complex cost-benefit analysis.[18] Agencies usually apply the standard of reasonable precautions on a source-specific basis. Typically, negligence-based regulations are enforced by iterative monitoring.[19]

[17] If source *A* is selected, the liability determination at that source ($10) will be applied to source *B*, for total liability of $20; if source *B* is selected, the $20 liability determination at that source will be applied to source *A*, for total liability of $40.

[18] Cf. W. Kip Viscusi et al., Economics of Regulation and Antitrust 35–37 (4th ed. 2005) (describing cost-benefit approaches to establishing regulatory norms and applying those approaches in various administrative contexts). For a helpful discussion of the cost-benefit analysis generally attendant to the negligence standard and its concomitant administrative costs, see Shavell, supra note 5, at 180–81, 185–89.

[19] For an example in the environmental regulation context, see infra note 40 and accompanying text.

As we explain below, SOS can be effectively employed to produce the same level of deterrence as IM under a variety of negligence-based regulatory regimes. First, we show how SOS applies to a negligence regulation that would lead the firm to take standardized compliance measures across all of its regulated sources. We then explain how SOS applies to a negligence regulation that would lead the firm to take source-specific or customized compliance measures.[20] Finally, we show how, with a minor modification, SOS applies to negligence-based regulation that evaluates all of the firm's risky activity against an aggregate standard of reasonableness.

1. Standardized Compliance

First, consider an IM-enforced negligence rule that would result in a firm taking the same compliance measures at all of its regulated sources. Assume that a firm operates two sources of risk, that the standard of care requires the firm to install a particular safety device costing $8 at each source, and that the agency threatens liability of $9 for each instance of noncompliance. The firm's aggregate expected liability and corresponding total expected cost under IM at various levels of compliance are set forth below:

Safety Devices Installed	*Compliance Cost*	*Aggregate Expected Liability*	*Total Expected Cost*[21]
2	$16	$0	$16
1	$8	$9	$17
0	$0	$18	$18

Under IM, the agency inspects both sources to determine the firm's liability at each—$0 for compliance or $9 for noncompliance—and therefore confronts the firm ex ante with aggregate liability of $0, $9,

[20] As we explain, SOS applies regardless of the nature of the precaution the firm would have adopted at each of its regulated sources under IM. See infra Subsection I.B.2.

[21] "Total expected cost" refers throughout to the sum of the costs of compliance and any liability imposed by the agency. We highlight the firm's aggregate expected liability for purposes of comparing IM and SOS, but the firm's incentives for compliance derive from the fact that compliance minimizes its total expected costs.

or $18 based on findings of full, partial, or zero compliance, respectively. Under SOS, the agency would randomly select and inspect only one of the regulated sources, determine liability for that source—either $0 for compliance or $9 for noncompliance—and automatically apply that liability outcome to all regulated sources. SOS thus confronts the firm with the same aggregate expected liability under each compliance alternative as the firm would face under IM: $0 for full compliance (a 100% probability of $0 in liability), $9 for partial compliance (a 50% probability of $18 in liability), and $18 for noncompliance (a 100% probability of $18 in liability).[22]

Because IM and SOS result in the same aggregate expected liability, both approaches motivate the firm to choose the same level of compliance—in this example, standardized compliance at both sources—because doing so minimizes its aggregate expected liability as well as its total expected costs.[23] Notably, SOS achieves this result without requiring the agency to determine the firm's liability at both sources.

2. Customized Compliance

A different negligence rule might result in the firm customizing compliance at one or more sources of risk. Assume, for example, that the firm in the above example is subject to a regulation requiring compliance measures costing $6 at source *A* and $10 at source *B* to prevent harm of $8 and $12 at each source, respectively. Assume also that the agency will impose liability of $8 if the firm fails to take precautions at source *A* and $12 if the firm fails to do so at source *B*. The firm's aggregate expected liability and correspond-

[22] Note that, in contrast to IM, under SOS ex post assessed liability may differ from ex ante expected liability for a given compliance strategy. Thus, in the example, SOS confronts the firm ex ante with aggregate expected liability of $18 for partial compliance, even though it is possible that the randomly selected source will be judged in compliance, and the agency will therefore impose $0 aggregate liability ex post.

[23] In the above example, SOS confronts a firm choosing noncompliance at one source with total expected costs of $17: expected liability of $9 and compliance costs of $8. A firm choosing noncompliance at both sources faces total expected costs of $18: expected liability of $18 and no additional compliance costs. Finally, if the firm selects full compliance, it expects total costs of $16: $16 in compliance costs and $0 expected liability. Thus, the firm minimizes total costs in this example by making the standardized investment necessary to comply at both sources.

ing total expected cost at each level of compliance under IM are described in the following table:

Source(s) in Compliance	*Compliance Costs*	*Liability*	*Total Cost*
A and B	$16	$0	$16
A only	$6	$12	$18
B only	$10	$8	$18
None	$0	$20	$20

Under IM, the firm eliminates its exposure to negligence liability by full compliance—taking the required customized precautions at both sources. Partial compliance results in liability of $12 or $8 for compliance only at source *A* or only at source *B*, respectively, and noncompliance at both sources results in liability of $20. Because the firm minimizes total expected cost to $16 by complying at both sources, IM leads the firm to full, customized compliance.

SOS confronts the firm with the same aggregate expected liability and total cost, thus inducing the same full, customized compliance. Under SOS, full compliance eliminates the firm's aggregate expected liability (the firm will face a 100% probability of $0 in liability). If the firm complies only at source *A*, it expects $12 in aggregate liability (a 50% probability of $24 in liability); if it complies only at source *B*, it expects $8 in liability (a 50% probability of $16 in liability); and if it complies at neither source, it expects $20 in liability (a 50% probability of $24 in liability and a 50% probability of $16 in liability). Including related compliance costs, the firm minimizes its total expected cost under SOS by complying with the regulation at both sources.[24]

Because SOS imposes greater liability for noncompliance at the randomly selected source than IM would for noncompliance at that

[24] If the firm complies at source *A* only, it expects liability of $12; including $6 in compliance costs, the firm expects total costs of $18. If the firm complies only at source *B*, it expects liability of $8; including $10 in compliance costs, the firm again expects total costs of $18. If the firm chooses total noncompliance, it faces expected liability—and total costs, as it pays nothing in compliance costs—of $20. Finally, if the firm complies with the regulation at both sources, it expects no liability; including the costs of compliance at both sources, it expects total costs of $16.

source, it might be thought that SOS will lead the firm to invest excessively in compliance at that source. But while SOS may impose greater liability than IM at a particular source and in the aggregate ex post, both SOS and IM generate the same expected liability at each source and in the aggregate ex ante—and hence produce the same incentives for compliance at each source and in the aggregate. For instance, if the firm in this example chooses to comply at source *A* only, under IM it bears expected liability of $12. The same is true under SOS. A firm complying only at source *A* faces the prospect that the agency will impose $24 in liability in the event that the agency randomly selects the noncompliant source *B*. This result does not lead the firm to take greater precautions at source *B*, however, because under SOS the firm internalizes that outcome with a 50% probability. Since it also internalizes the 50% chance that the agency will select the compliant source *A* and impose $0 in liability, the firm anticipates aggregate liability of $12 for its failure to comply at source *B*. More generally, because the firm will follow the compliance strategy that minimizes its ex ante total costs—and because ex ante total costs are the same for each compliance strategy under IM and SOS—the firm will choose the same level of compliance under SOS as under IM, even though SOS may result in greater ex post liability at an individual source and in the aggregate.

3. Aggregate Liability

Aggregate liability refers to regulations in which an agency ex ante specifies the reasonable level of risk collectively for all regulated sources and then assesses liability by comparing the total actual risk output to the total permitted level.[25] Aggregate liability regulation is generally thought to reduce the costs of compliance because the firm itself, presumably having superior information, selects the most efficient means of meeting the aggregate risk standard for

[25] The aggregate liability approach is essentially a variation on the negligence rule. We address this approach separately because some agencies, particularly the EPA, consider it an efficient alternative to source-specific, negligence-based regulation. The approach is often referred to as "bubble" treatment of multiple sources; the term is commonly associated with EPA regulations in the emissions context. See Chevron U.S.A. v. Natural Res. Def. Council, 467 U.S. 837, 855 (1984).

all sources covered by the regulation.[26] While reducing the costs of *compliance* with regulatory standards, however, aggregate liability does not reduce the costs of *monitoring* whether a firm is in compliance if regulators use IM to determine aggregate compliance levels. IM is currently the standard operating procedure for enforcing aggregate liability regulations, and agencies therefore incur the costly step of measuring risk at each source simply to derive the total actual output for comparison with the total permitted output.[27]

SOS can effectively replace IM in this context. Under SOS, regulators would randomly select a source, determine risk at that source, apply that result to all other regulated sources controlled by the firm, and impose liability where total risk exceeds the aggregate limit. For illustration, consider a firm with two factories, *A* and *B*, each producing a regulated chemical. Assume that the agency sets the aggregate limit at 16 units and that the regulator imposes liability of $1 for each unit in excess of that limit. Finally, assume that the firm has three alternative compliance strategies with correspondent chemical output and expected liability. These factors and the firm's expected costs are summarized in the following table:

Compliance Alternative	*Investment in Precautions*		*Chemical Production (Units)*		*Expected Liability*	*Total Costs*
	Source A	*Source B*	*Source A*	*Source B*		
1	$7	$7	8	8	$0	$14
2	$7	$0	8	25	$17	$24
3	$0	$0	25	25	$34	$34

Under IM, the firm faces aggregate expected liability of $0 if it chooses full compliance, $17 if it complies at one source, and $34 if it does not comply at all. The firm therefore minimizes expected li-

[26] The firm need not comply with agency-specified technology and risk limits at individual sources; instead, to avoid liability, it must only ensure that its total risk at all regulated sources does not exceed the total risk limit. The firm will therefore seek the lowest-cost adjustments to levels of activity and care at the regulated sources necessary to comply with the aggregate limit. See generally Tom Tietenberg, Environmental and Natural Resource Economics 365–66 (4th ed. 1996); T.H. Tietenberg, Emissions Trading: An Exercise in Reforming Pollution Policy 51–52 (1985).

[27] The aggregate liability approach does not require monitoring at each source; the agency could instead assess total risk and compare it to the aggregate benchmark.

ability and total cost by complying—that is, by investing $7 in precautions—at both sources.

The firm faces identical aggregate expected liability, and identical incentives for compliance, in each scenario under SOS. If the firm complies fully, it faces expected liability of $0 (a 100% probability that the agency will impose liability of $0); if it complies at one source, it expects liability of $17 (a 50% probability that the agency will impose liability of $34); and if it chooses total noncompliance, it expects $34 in liability (a 100% probability that the agency will impose liability of $34). Under SOS, the firm minimizes its aggregate liability and hence its total costs by complying fully.[28] Thus, as under IM, a firm facing SOS would be led to select full compliance.

SOS may not replicate deterrence under IM, however, in cases of partial compliance where its application could produce a finding that the firm's aggregate output falls *below the aggregate limit set by the agency and therefore warrants zero liability*.[29] In such a case, SOS will result in greater aggregate expected liability than IM. For illustration, consider the firm in the previous example, but assume that by taking precautions the firm will reduce production of the regulated chemical to 0, while the aggregate limit remains 16 units. The firm's expected costs under each compliance scenario are summarized in the following table:

Compliance Alternative	*Investment in Precautions*		*Chemical Production (Units)*		*Expected Liability*	*Total Costs*
	Source A	*Source B*	*Source A*	*Source B*		
1	$7	$7	0	0	$0	$14
2	$7	$0	0	25	$9	$16
3	$0	$0	25	25	$34	$34

[28] The firm's total costs under SOS are identical under each compliance alternative to the firm's total costs under IM. If the firm chooses full compliance, its total costs are $14, including $14 in compliance costs and $0 in expected liability. If the firm chooses partial compliance, its total costs are $24, including $7 in compliance costs and $17 in expected liability. Finally, if the firm chooses total noncompliance, its total costs are $34, including $0 in compliance costs and $34 in aggregate expected liability.

[29] Note that, in the previous example, the firm's risk output from all regulated sources equaled or exceeded the aggregate output limit imposed by the agency in each compliance scenario.

Consider the firm's expected liability when it complies at only one source—that is, where the firm invests $7 at one source and $0 at the other. Under IM, the firm faces aggregate expected liability of $9 (the firm has total output of 25 units, 9 in excess of the agency's aggregate limit of 16). Under SOS, however, the firm faces $17 in expected liability if it complies at one source: a 50% probability that the agency will impose $34 in liability if it selects the noncompliant source. In some cases, the increase in aggregate expected liability under SOS could distort incentives relative to IM. The greater threat of liability under SOS may, for example, induce the firm to overinvest in precautions given its uncertainty regarding what aggregate limit the agency will impose on output or its anticipation of an erroneously stringent limit.[30]

A minor modification to the SOS procedure will permit SOS to duplicate aggregate expected liability in all cases, regardless whether the measured risk output exceeds or falls below the aggregate limit set by the agency. The modification requires the firm to specify, at some point before the agency randomly selects a source for monitoring, the expected or actual output at each of its regulated sources. Under the modified approach, the agency calculates the firm's aggregate excess output as the sum of two findings: first, the excess of total specified output over total permitted output; and second, based on assessment of a randomly selected source, the excess of actual output over specified output applied to all regulated sources. Multiplying this aggregate excess output by the prescribed sanction yields the firm's aggregate liability.[31]

To illustrate how SOS with specification replicates aggregate expected liability, consider again the previous example. Under IM, the firm expects aggregate liability of $9 when complying only at one source. Now consider SOS with specification, and suppose that

[30] See Shavell, supra note 5, at 224–28, 253 n.36 (discussing the potential for the negligence rule to distort incentives for compliance); David Rosenberg, The Causal Connection in Mass Exposure Cases: A "Public Law" Vision of the Tort System, 97 Harv. L. Rev. 849, 863–65 (1984) (explaining how magnified negligence liability can induce socially excessive investment in care).

[31] Alternatively, the agency could apply SOS without specification, but when it results in a total risk determination less than the cumulative standard set by the agency, regulators would credit the difference to the firm. Whether this approach or our specification modification would be less expensive for firms and regulators is an empirical question beyond the scope of this Essay.

the firm specifies actual output of 0 and 25 at sources *A* and *B*, respectively. First, the agency's comparison of the firm's total specified output to total permitted output would result in a finding of 9 excess units (25 specified minus 16 permitted). Second, assuming that the firm's specifications accurately reflect its actual outputs, the agency's random inspection of one source would reveal no excess output. Consequently, the firm's total excess of 9 units would translate into aggregate liability of $9, the same sanction the firm faces under IM. Now suppose, however, that the firm's actual output exceeds its specifications; for example, assume that the actual output at source *A* is 25 rather than the specified output of 0. (For simplicity, assume that output at *B* is accurately specified at 25.) Under IM, the firm would bear aggregate expected liability of $34: 50 actual units minus 16 permitted. SOS with specification also confronts the firm with aggregate expected liability of $34: the agency will impose $9 in liability for the 9-unit excess in total specified output, and the firm will face a 50% chance that the agency will randomly select source *A*, apply its finding of 25 units over the specified output to both sources, and impose $50 in additional liability.[32]

SOS with specification will not impose more costs on firms than IM because firms must determine risk output at each source in any event to plan and monitor their compliance with the aggregate risk limit set by the agency.[33] Thus, agencies will not need to sort out

[32] Firms could not profit from underspecifying the level of actual risk at any source. Any gains from such contrivance would be fully offset by the corresponding increase in the expected liability the firm faces for excess output at the randomly selected source. Consider, for example, a firm with actual output of 25 units at each of two sources, and suppose, in an attempt to appear in compliance with the permitted total output of 16 units, the firm specifies 8 units of output at each source. Although the agency would impose no liability for excess total output, expected liability from the random selection of either source would expose the firm to the proper amount of aggregate expected liability of $34. (The firm would face a 100% probability that the agency would find the selected source 17 units in excess of its specified output, attribute that result to both sources, and impose liability of $34.) Note that IM would generate the same aggregate expected liability—the agency would detect excess output of 17 units at each source—but require the greater expense of monitoring both sources instead of just one.

[33] In practice, agencies employing aggregate liability regulations generally require firms to disclose risk outputs at each regulated source as a condition of agency authorization to engage in the risky activity. For a discussion of this point in the EPA context, see infra note 57 and source cited therein.

the cases in which specification would be necessary. Using SOS with specification in every case will allow agencies enforcing aggregate liability regulations to replicate firms' aggregate expected liability at much lower cost than the current IM regime.

C. Frequency of Liability Determination

In addition to establishing the method for determining liability and the magnitude of the sanction, regulators in many contexts choose the frequency with which liability will be assessed. SOS can also be used to replicate the firm's expected liability over a series of monitoring periods while reducing the frequency of monitoring.[34] To apply SOS in this context, regulators would determine the firm's liability at a particular point in time and apply that determination perforce to all other time periods during which the firm would otherwise be subject to IM.

To see this, consider a firm operating a single source of risk and assume that the level of risk at that source varies over time: at $t = 1$, the source creates 10 units of risk, while at $t = 2$, the source creates 20 units. Assume, too, that regulators will impose liability of $1 for each unit of risk at each point in time.[35] The following table describes the firm's expected liability under iterative monitoring:

Monitoring Period	*Level of Risk*	*Expected Liability*
1	10	$10
2	20	$20
Total	**30**	**$30**

Under IM, the firm faces aggregate expected liability of $30: $10 assessed at $t = 1$ and $20 assessed at $t = 2$. Under SOS, rather than

[34] Of course, SOS can also replace an IM regime operating on multiple dimensions—for example, a regulatory approach requiring monitoring of a set of sources periodically and monitoring of each source within the set individually during each period.

[35] For ease of illustration, we assume that the regulator imposes liability strictly based on the level of risk during each monitoring period. We note, however, that SOS could also be used to reduce the frequency of negligence-based determinations of liability.

monitoring at both $t = 1$ and $t = 2$, regulators will randomly select one period, determine the firm's liability during that period, and apply that determination to both periods. If the agency monitors at $t = 1$, it will impose liability of $10 for that period; applying that result to $t = 2$, it will impose total liability of $20. If the agency monitors at $t = 2$, it will impose $20 in liability for that period and thus impose total liability of $40. Anticipating each outcome with 50% probability, the firm's aggregate expected liability under SOS is $30—precisely the same as under IM.

The application of SOS to periodic monitoring will be straightforward when the agency can randomly select any period in the relevant series—first, last, or any interim period—and impose liability without distorting the firm's incentives for compliance. If, however, the agency randomly selects any time period before the last, and the firm is able to alter its compliance strategy during ensuing periods, there is some risk that the firm will lack incentives for compliance thereafter because its liability during future periods has been determined by monitoring that has already taken place.[36]

An agency can address this problem in two ways. First, where reliable records indicate risk output during all relevant periods, the agency can wait until the last period in the series and then randomly select any period for monitoring and apply the liability determination to all periods. Where such records (or their functional equivalents) are not available, the agency could apply the liability outcome during a randomly selected period only to preceding periods of time. Alternatively, the agency could set a given probability of monitoring for each period and multiply any liability imposed during selected periods by the reciprocal of the probability that the period will be selected.[37]

Whether agencies should assess liability at individual sources or in the aggregate; whether they should employ a negligence-based or strict liability approach to determining liability; or whether

[36] Of course, for incentives for compliance to be compromised, the firm must also be *aware* that the agency has selected a previous period as determinative of liability. In many regulatory contexts, the firm may not know or may be kept unaware of the period of time selected for monitoring.

[37] Cf. David Rosenberg & Steven Shavell, A Simple Proposal to Halve Litigation Costs, 91 Va. L. Rev. 1721, 1721 (2005) (describing a system in which half of civil cases would be dismissed and damages would be doubled in remaining cases).

monitoring should be calibrated in some combination of sources, time periods, or otherwise are matters beyond the scope of this Essay. Our point is simply that SOS can replicate the deterrent effect of IM regardless of how the agency might resolve those questions.

II. BENEFITS OF SOS: THE CASE OF THE EPA

The principal benefit of SOS is a significant reduction in the cost of monitoring without affecting incentives for compliance. SOS generates substantial cost savings for a straightforward reason: under SOS, the regulator monitors only *one* regulated source, while under the corresponding IM regime it must monitor *every* regulated source operated by the firm. The cost savings generated by SOS increase in a linear fashion as the number of sources controlled by each firm increases. These savings may yield additional enforcement gains if regulators redeploy resources freed up by SOS or otherwise exploit its efficiencies to pursue previously underserved regulatory priorities.

We expect that agencies employing SOS would significantly reduce the monitoring costs they face under IM. In this Part, we briefly examine the savings SOS would generate if the EPA used it to regulate emissions of airborne pollutants from "stationary sources" such as industrial smokestacks.[38] In general, the EPA specifies the type of precautions required at each source based on a negligence-type cost-benefit analysis[39] and ensures compliance with these precautions by iteratively monitoring each regulated source.[40]

[38] 42 U.S.C. § 7502(c) (2000). In practice, most monitoring of compliance with this mandate is conducted by state officials, who propose, promulgate, and enforce state implementation plans pursuant to the Clean Air Act. For convenience, however, we refer jointly to the state and federal authorities charged with monitoring compliance with the Clean Air Act as the "EPA."

[39] See, e.g., Chevron U.S.A. v. Natural Res. Def. Council, 467 U.S. 837, 847 (1984); Tietenberg, Environmental and Natural Resource Economics, supra note 26, at 83–85.

[40] EPA and state officials engage in site visits that permit direct monitoring of compliance at stationary sources. See EPA, Clean Air Act Stationary Source Compliance Monitoring Strategy 7–8 (2001), available at http://www.epa.gov/compliance/resources/policies/monitoring/cmspolicy.pdf (describing the type and frequency of visits recommended by the EPA). Where a violation of relevant emissions standards is detected, the EPA imposes liability following guidelines that measure, inter alia, the actual or possible harm, the toxicity of the pollutant, and the size of the firm. See Memorandum from the EPA Office of Enforcement and Compliance Assurance on Clarification of the Use of Appendix I of the Clean Air

Monitoring every potential source of airborne pollution in the United States, however, would be prohibitively expensive. As a result, Congress has limited EPA oversight to "major" sources, defined as "any stationary facility or source of air pollutants which directly emits, or has the potential to emit, one hundred tons per year of any air pollutant."[41] There are approximately 22,000 such sources in the United States.[42] On average, a single firm operates approximately four major stationary sources,[43] although the distribution of sources among firms likely varies by industry, geography, and scale of the firm.

The burden of iteratively monitoring thousands of major sources has led the EPA to reduce the frequency of such monitoring. Currently, the agency aims to test each source at least once every five years.[44] Despite this reduced frequency, EPA's monitoring obligations consume more than $130 million per year[45]—

Act Stationary Source Civil Penalty Policy (July 23, 1994), available at http://www.epa.gov/compliance/resources/policies/civil/caa/stationary/penpol.pdf; see also David B. Spence, The Shadow of the Rational Polluter: Rethinking the Role of Rational Actor Models in Environmental Law, 89 Cal. L. Rev. 917, 937 (2001) (describing EPA procedures for imposing liability).

[41] 42 U.S.C. § 7602(j); see also 40 C.F.R. § 52.21(b)(1)(i) (2007) (defining "major stationary source").

[42] EPA, Air Pollution Operating Permit Program Update 2 (1998).

[43] See EPA, Regulatory Impact Analysis for Part 64 Compliance Assurance Monitoring Regulation 23–24 (1997) (on file with the Virginia Law Review Association) (stating that, in a representative sample of five states, the EPA identified 1179 facilities containing 4642 major stationary sources).

[44] See EPA, supra note 40, at 7. In order to manage monitoring costs, the EPA also uses a number of statistical thresholds to dictate whether additional cost-intensive monitoring is necessary. For example, the EPA recommends that state investigators conduct a "full compliance evaluation" of any facility with a major stationary source within 20% of the maximum emissions authorized by regulatory standards. Id. Moreover, to reduce monitoring costs the EPA reviews compliance materials provided by the firms themselves, relying on those estimates for "reasonable assurance of compliance with emission limitations or standards for the anticipated range of operations at a pollutant-specific emissions unit." 40 C.F.R. § 64.3(a) (2007).

[45] We assume that the average expense of directly monitoring a stationary source is $30,000. See EPA Office of the Inspector General, Report of EPA's Oversight of State Stack Testing Programs 2 (2000), available at http://www.epa.gov/oig/reports/2000/stack.pdf (estimating the costs of such monitoring at $10,000 to $50,000 per source). We also assume that 22,000 sources are subject to such monitoring and that each source is monitored once every five years. See supra text accompanying note 44. EPA and state regulators pass these costs on to regulated firms, which are statutorily required to pay for the reasonable costs of enforcing permits issued under the Clean Air Act. 42 U.S.C. § 7410(a)(2)(L)(ii). Because firms

and the agency has failed to fulfill even these commitments on time.[46]

It is apparent, then, that the EPA could reap substantial benefits by replacing its iterative monitoring of stationary sources with single-outcome sampling. Indeed, given the EPA's estimated average of four sources per firm,[47] SOS would reduce monitoring costs by 75%, or approximately $98 million per year. The precise magnitude of the savings, of course, may be greater depending on whether larger-scale firms operate more sources than the national average. In any event, these savings should prove significant for an agency frequently described as "desperately short" of enforcement resources.[48] SOS would therefore free the EPA and firms of the unnecessary costs of IM and enable the agency to use enforcement strategies that are currently too costly to pursue.[49]

SOS may also permit the EPA to make greater use of strict liability or aggregate liability regulations. Notwithstanding the apparent efficiency of these modes of regulation in comparison to negligence-based, source-specific regulation,[50] Congress and the

must also prepare to review regulators' on-site inspections and their findings, the total costs of iterative monitoring of emissions compliance may well be higher than those included in our estimate.

[46] According to a recent study, only 14% of major sources were tested even once within a ten year period. See EPA Office of the Inspector General, supra note 45, at 10 (drawing this conclusion from an analysis of the EPA's nationwide stack-testing database); see also EPA Office of the Inspector General, Region 6's Oversight of New Mexico Enforcement Data 7 (1998), available at http://epa.gov/oig/reports/1998/8100078.pdf.

[47] See supra note 43 and accompanying text.

[48] Rena I. Steinzor, Devolution and the Public Health, 24 Harv. Envtl. L. Rev. 351, 354 (2000) ("The simple truth is that EPA and the states are desperately short of money, with statutory mandates that far exceed their administrative and enforcement capabilities.").

[49] Although this Part emphasizes the number of major stationary sources subject to on-site inspection, we note that SOS could also be applied to reduce the frequency of EPA inspections without diminishing current levels of deterrence. See supra Section I.C.

[50] Several commentators have argued that the cost-benefit, command-and-control regulatory structure of most EPA regulations is administratively burdensome. See, e.g., Shi-Ling Hsu, Fairness Versus Efficiency in Environmental Law, 31 Ecology L.Q. 303, 379 (2004) (describing the administration of the EPA's "traditional 'command-and-control' regulatory structure" as "difficult, time-consuming, and costly" and noting that the EPA is responsible for setting discharge standards for more than twenty-seven different types of pollution sources under one provision of the Clean Water Act alone).

EPA have limited use of strict liability[51] and aggregate liability (or "bubble") regulation.[52] Commentators have noted that significant monitoring costs attendant to such regulations may have caused the EPA to refrain from making greater use of strict liability and aggregate liability regulations.[53] By reducing the monitoring costs associated with strict liability and aggregate liability regulations, SOS may permit the EPA and lawmakers to consider wider use of these enforcement approaches.

In particular, the EPA has long resisted calls for expanded use of strict liability in the form of taxes calculated on the basis of units of pollution or risk.[54] One objection has been that the benefits of such a regime would be outweighed by its administrative costs, including the cost of iteratively monitoring to assess liability at every source

[51] Despite analysis supporting wider use of strict liability, see, e.g., Robert A. Pulver, Liability Rules as a Solution to the Problem of Waste in Western Water Law: An Economic Analysis, 76 Cal. L. Rev. 671, 707 n.153 (1988), Congress has limited this approach to a small set of cases involving highly hazardous substances, see Comprehensive Environmental Response, Compensation, and Liability Act ("CERCLA"), 42 U.S.C. §§ 9601–9609 (2000). EPA regulations are almost exclusively "standard-setting," like those under the Clean Air Act; CERCLA is an exception to this general rule. See, e.g., New York v. Shore Realty Corp., 759 F.2d 1032, 1041–42 (2d Cir. 1985) ("CERCLA is not a regulatory standard-setting statute such as the Clean Air Act. . . . Congress intended that responsible parties be held strictly liable . . . .").

[52] The EPA applies aggregate liability standards in its "bubble" treatment of multiple stationary sources of emissions. Chevron U.S.A. v. Natural Res. Def. Council, 467 U.S. 837, 848–53 (1984); Robert W. Hahn & Gordon L. Hester, Where Did All the Markets Go? An Analysis of EPA's Emissions Trading Program, 6 Yale J. on Reg. 109, 123 (1989). Under this approach, "[i]nstead of trying to meet the emission limit set for a specific [source], a plant is allowed to arrange its emission controls such that the facility satisfies its aggregate permitted amount for a specific pollutant." Kathryn C. Wilson, The International Air Quality Management District: Is Emissions Trading the Innovative Solution to the Transboundary Pollution Problem?, 30 Tex. Int'l L.J. 369, 381 (1995). Bubble regulations, however, apply to just 1% of the stationary sources in the United States. Thus, notwithstanding the significant potential benefits of "bubble" treatment of multiple sources, several commentators have concluded that "the net impact of bubbles on environmental quality has not been significant." Hahn & Hester, supra, at 129 & n.105 (describing internal EPA analysis noting that "bubble" treatment of multiple stationary sources has generated surprisingly small benefits).

[53] See, e.g., Daniel H. Cole, Pollution and Property 82 (2002).

[54] In the administrative context, such liability is commonly modeled on Pigouvian taxes, which charge the firm for the expected social cost of risk rather than actualized harm. See A.C. Pigou, Wealth and Welfare 164 (1912); see also Hsu, supra note 50, at 400 (proposing, but recognizing the "political realities" counseling against, the use of Pigouvian taxation by the EPA).

of risk rather than only sources presenting readily identifiable breaches of a prescribed negligence-based standard.[55] Application of SOS to strict liability taxes will, as we have shown, maintain deterrence at the level achieved by IM but reduce monitoring cost by orders of magnitude. This should render such taxes a more palatable regulatory alternative.

The EPA's use of aggregate liability, or "bubble," regulation also appears to have been curtailed by the significant costs of iterative monitoring.[56] Use of SOS would permit the EPA to reduce these costs significantly. Moreover, under current law, firms seeking bubble treatment of multiple sources must disclose target emissions at each source;[57] thus, the specification requirement needed for effective implementation of SOS would impose no additional costs on regulators or firms.[58] Switching to SOS would reduce the costs of monitoring compliance with bubble-based emissions limits without altering firms' incentives for compliance with those limits. This reduction in cost should make bubble regulation far more attractive as an enforcement strategy.

In sum, the EPA's use of SOS rather than iterative monitoring would produce three significant benefits. First, SOS would reduce the costs of monitoring compliance under the predominant strategy of negligence-type, individual-source regulation. Second, SOS may enable the EPA to expand its use of alternative modes of regulation, including strict taxes and bubble-based regulation. Finally, the savings from SOS should free resources for reallocation to more efficient uses and render underused but potentially productive strategies more cost effective.

---

[55] E.g., David M. Driesen, The Economic Dynamics of Environmental Law 69 (2003) (expressing the view that the EPA would be administratively incapable of implementing a Pigouvian tax).

[56] Cole, supra note 53, at 82 (noting that the success of bubble regulation of emissions has been limited by the cost of adequately monitoring individual source emissions).

[57] Before regulators approve a group of sources for bubble treatment, the firm is required to disclose the means of compliance with the bubble's aggregate limit. See 40 C.F.R. § 52.21(aa)(3)(i-iii) (2007) (requiring firms seeking bubble treatment to specify each source within the proposed bubble and expected actual emissions at each source).

[58] See supra text accompanying notes 31–33 for a discussion of the specification requirement used to apply SOS to aggregate liability regulations.

III. POTENTIAL COSTS OF SOS

In this Part, we address risk-bearing and judgment-proof costs associated with the use of SOS. Both costs arise from the possibility that SOS may in some cases produce highly variable liability outcomes, which might include the threat of a result imposing so much total liability as to consume much or even all of the firm's wealth. First, we consider the case in which the prospect of large total liability relative to the firm's assets may burden a risk-averse firm with increased risk-bearing costs.[59] Second, we examine the case in which total liability might exceed the firm's assets and, anticipating that it could be judgment-proof, the firm may have less incentive for compliance.

The potential for such problems is context specific—depending, for example, on the firm's degree of risk aversion and its wealth relative to expected liability.[60] Of course, the ultimate test of the utility of our proposal is not whether these problems will arise but whether they raise the total social costs of SOS above those of the corresponding IM regime. There is good reason to believe that SOS will generally pass this test because, as we explain, risk-bearing and judgment-proof problems are both unlikely to impose substantial costs and are amenable to the same solution: wider sampling paid for by the firm.

A. Risk-Bearing Cost

The prospect that SOS might result in the imposition of significant total liability might increase the firm's risk-bearing costs and consequently distort its incentives for compliance. For two reasons, we think these costs are unlikely to be substantial. First, firms are

[59] We think it is unlikely that the agency would be risk-averse with respect to the opposite outcome, that is, that no liability will be imposed as a result of random selection of a single source. The mechanical nature of SOS would likely shield the administrator from any political costs associated with such a result.

[60] The magnitude of such problems will also depend on the nature and structure of the regulatory regime and the degree to which liability outcomes may vary under SOS. For example, if the magnitude of a given sanction was set in part to recoup the costs of IM, the cost savings from using SOS might lead the agency to lower the sanction, reducing the variability of liability outcomes under SOS. Settlement or other negotiated arrangements may also mitigate or eliminate these problems.

generally structured to minimize the impact of risk on their assets.[61] Second, insurance for regulatory liability is generally available.[62]

Even if SOS poses a danger of increased risk-bearing costs in some cases, the agency could readily address the problem by providing firms with the option of having liability determined on the basis of wider sampling. This option would be exercisable by the firm, which could choose to have the agency randomly select and determine liability at any number of its regulated sources. In the event that the firm chooses to have the agency conduct wider sampling, the average of all liability determinations at the selected sources will automatically determine the firm's liability at all of its regulated sources.[63]

Two conditions should govern the availability of this option. First, the option should be made available only *before* the regulator has randomly selected any source for monitoring. If the firm is permitted to elect wider sampling after it has learned what source the agency has randomly selected, the firm might opt for wider sampling opportunistically, diluting deterrence.[64]

Second, the firm must pay the cost incurred by the agency in monitoring additional sources. This condition prevents the firm from unnecessarily exercising the option. Charging firms for addi-

[61] See David Mayers & Clifford W. Smith, Jr., On the Corporate Demand for Insurance, 55 J. Bus. 281, 281–82 (1982).

[62] In the environmental context, such insurance is generally available for regulatory liability incurred accidentally. See EPA, Environmental Insurance Policy Coverage and Terms 1, http://www.epa.gov/brownfields/insurance/policy_coverage_chart.pdf (last visited Sept. 17, 2007) (stating that "most" insurers offer policies covering "[d]amages (including pollution liability) from acts, errors, or omissions" related to compliance); see also D. Evan Van Hook et al., The Challenge of Brownfield Clusters: Implementing a Multi-Site Approach for Brownfield Remediation and Reuse, 12 N.Y.U. Envtl. L.J. 111, 149 n.122 (2003) (describing the "widely recognized" value of environmental insurance). Although the maximum coverage under such a policy would have to be higher for a firm facing SOS than under IM, the premiums paid for such coverage should not rise: the prospects of higher and lower liability outcomes offset each other.

[63] We refer here to "wider sampling" in terms of the SOS model of random selection within a given sample, as distinguished from *statistical* sampling, which requires the regulator to determine the size and composition of the sample. See supra note 6.

[64] Specifically, the firm would request wider sampling only when the initially selected source would generate above-average total liability and would not request wider sampling when the initially selected source would generate below-average total liability. Using this strategy, the firm would be confronted ex ante with lower aggregate expected liability than it would face under the corresponding IM regime.

tional sampling also relieves the agency of the costs of determining whether wider sampling is needed and promotes deterrence by requiring the firm to internalize enforcement costs associated with its risky activity.[65] In short, the firm will exercise the option for wider sampling only when doing so is socially appropriate.

To illustrate the value of this option to a risk-averse firm, consider a firm with sources *A*, *B*, and *C*. Assume that the sources produce risk of 10, 20, and 30 units at each source respectively and that the governing agency will assess liability of $1 per unit of risk.[66] Finally, assume that the firm starts with wealth of $100 and that it is risk averse and therefore attaches diminishing marginal utility to money.

Under IM, the firm expects aggregate liability of $60.[67] Under SOS, of course, the firm's aggregate expected liability is identical: with probability of 33%, it expects liability of $30, $60, or $90 if source *A*, *B*, or *C* is selected, respectively.[68] Thus, the firm's aggregate expected liability under SOS is also $60.[69] Under SOS, however, the firm is exposed to the possibility that the agency will impose $90 in liability, leaving it with just $10 in wealth.[70] Suppose

[65] Of course, to promote full internalization of the social costs of risky activity, the firm may also be required to pay the costs of monitoring even a single source under SOS. See Shavell, supra note 5, at 178, 411–12 (discussing the social optimality of internalization of enforcement costs). Firms subject to EPA regulation of stationary sources of airborne emissions are required to pay such costs. See supra note 45.

[66] For simplicity we assume here that the agency imposes liability strictly; but, as we have explained, application of SOS extends to negligence-based regulation as well. See supra Section I.B.

[67] Using the standard practice of modeling the diminishing marginal utility of money by equating the welfare derived from a given amount of money with the square root of that amount, we can calculate the firm's expected welfare under IM. Anticipating with certainty that the agency will monitor all three sources and assess liability of $60, the firm expects that its total wealth will be reduced to $40, and thus welfare under IM will be equal to 6.32, as compared to pre-liability wealth of $100 and utility of 10.

[68] Applying the liability determination at the selected source to all of the firm's sources, including the selected source and those not selected, a liability determination of $10 at source *A* results in aggregate liability of $30; a liability determination of $20 at source *B* results in aggregate liability of $60; and a liability determination of $30 at source *C* results in aggregate liability of $90.

[69] This is simply the sum of (33% x $30), (33% x $60), and (33% x $90).

[70] Because the firm anticipates aggregate liability of $60 under IM, we would expect that the other possible liability outcomes under SOS—$30 or $60—would not cause the firm to incur any SOS-specific risk-bearing costs.

that the possibility of the $90 total liability outcome increases the firm's risk-bearing costs above those it would bear under IM.[71]

Now assume that such risk-bearing costs exceed the costs of wider sampling, so that the firm would exercise its option to have the agency randomly select two sources rather than one. Despite exercising the option, the firm faces the same aggregate expected liability of $60.[72] However, the possible liability outcomes of $45, $60, and $75 no longer present the firm with the risk that it will be exposed to $90 in liability, eliminating the risk-bearing costs associated with that possibility.[73]

Whether this firm will choose to have one, two, or three sources randomly sampled will depend, of course, on the relationship between its utility under each alternative and the cost of additional

[71] If the agency selects source *C* and imposes liability of $90, this would reduce firm wealth to $10 and welfare to 3.16. If the agency selects source *B* and imposes liability of $60, this would reduce firm wealth to $40 and welfare to 6.32. Finally, if the agency selects source *A* and imposes liability of $30, this would reduce firm wealth to $70 and welfare to 8.37. Expecting each state of the world with 33% probability, the firm's total expected welfare under SOS is 5.95, lower than the firm's expected welfare of 6.32 when IM is used to impose the same aggregate expected liability. See supra note 67.

[72] If the agency selects sources *A* and *B*, assessing liability of $10 and $20, respectively, and applying the $15 average liability determination to all three sources, it will impose $45 in total liability. If the agency selects sources *B* and *C*, assessing liability of $20 and $30, respectively, and applying the $25 average liability determination to all three sources, it will impose $75 in total liability. Finally, if the agency selects sources *A* and *C*, assessing liability of $10 and $30, respectively, and applying the $20 average to all three sources, it will impose total liability of $60. The firm expects each outcome with 33% probability and thus aggregate expected liability of $60, or the sum of (33% x $45), (33% x $75), and (33% x $60).

[73] To see this in terms of the utility calculus described above in note 71, we can calculate the firm's utility under SOS with two sources sampled. If the firm exercises its option to have two sources sampled, it will be exposed either to liability of $45, reducing wealth to $55 and welfare to 7.42; liability of $75, reducing wealth to $25 and welfare to 5.00; or liability of $60, reducing wealth to $40 and welfare to 6.32. The firm expects each outcome with 33% probability and thus its expected welfare is 6.25, or the sum of (33% x 7.42), (33% x 5.00), and (33% x 6.32). Note that the firm's expected welfare is higher under SOS when two sources are sampled than when a single source is sampled, see supra note 71, because the sampling of an additional source decreases the variability of the firm's aggregate expected liability. Of course, if the firm also incurred SOS-specific risk-bearing costs from the prospect of $75 in liability—still a possibility under SOS where the firm exercises its option to have two sources selected—the firm could elect to have all three sources sampled to ensure that its aggregate liability would be exactly $60.

sampling.[74] By providing the firm with the option to choose the number of sources sampled, however, the agency ensures that the firm internalizes the costs of its risk-producing activity without imposing any additional informational or decisional costs on regulators.[75]

B. Judgment-Proof Cost

SOS is susceptible to the judgment-proof problem in a case where total liability under SOS exceeds the firm's assets and liability insurance. This possibility of excess liability can compromise deterrence. If the firm anticipates being unable to pay certain high-liability outcomes under SOS, it may reduce investments in regulatory compliance.[76]

[74] In terms of the utility calculus described above, see supra notes 67–73, the firm will compare the marginal costs of the additional sampling with the increase in expected welfare provided by such sampling to determine whether to exercise its option to have two sources sampled. In this case, the firm can pay as much as $3.50 for sampling of the second source and still expect total welfare higher than its welfare when a single source is sampled. If the cost of the second sample were $3.60, however, the firm's expected welfare would fall to 5.94, and the firm would prefer sampling of a single source. See supra note 71. While this choice to forgo wider sampling would result in greater risk-bearing costs for the firm, because the costs of additional sampling exceed these risk-bearing costs, SOS remains the socially superior alternative to IM, which definitionally requires sampling of all regulated sources.

[75] Regulators could also address risk-bearing cost by calculating aggregate liability based on a specification procedure similar to the one we propose for use in the aggregate liability context. See supra Subsection I.B.3. Under this alternative procedure, the firm would specify ex ante the proportion of total risk produced by each regulated source. The agency would then randomly select a source, weighting the probability of selecting each source by the specified proportion; determine liability at the selected source; and impose aggregate liability equal to liability at the selected source divided by the specified proportion for that source. This procedure has the advantage of reducing risk-bearing cost without requiring additional sampling. However, the procedure is also more susceptible to deliberate judgment-proofing. Under SOS, aggregate liability is determined by the number of regulated sources controlled by the firm, limiting the variability in potential outcomes. The proportion-specifying procedure, however, would permit the firm deliberately to expose itself to liability much greater than its assets, with no limit on the variability among potential outcomes. Whether the potential costs of the specification procedure described here would outweigh its benefits in a particular regulatory context is an empirical question beyond the scope of this Essay. We thank Michael Abramowicz for calling this use of the specification procedure to our attention.

[76] We distinguish between this type of judgment-proof problem, arising because of variance in liability outcomes, see Rohan Pitchford, How Liable Should a Lender Be? The Case of Judgment-Proof Firms and Environmental Risk, 85 Am. Econ. Rev. 1171,

This problem is unlikely to arise.[77] Agencies may minimize the likelihood that a firm will be judgment-proof by determining liability on the basis of risk rather than actualized harm.[78] Moreover, small firms—those most likely to have limited assets relative to aggregate expected liability—will likely have few regulated sources, reducing the likelihood that SOS will result in the imposition of high total liability.[79] Further, the problem may be mitigated in cases where the frequency of agency inspections effectively reduces the number of sources subject to monitoring at one time.[80] Regulators may also have authority to institute safeguards against the judgment-proof problem, for example through minimum-asset and li-

1171–72 (1995), and the type arising because firms deliberately judgment-proof themselves, for example by encumbering assets with secured debt, see Lynn M. LoPucki, The Death of Liability, 106 Yale L.J. 1, 14–15 (1996). Here we address only the former—that is, the judgment-proof problems created by the variance in liability outcomes generated by the SOS mechanism. For comprehensive consideration of the problems presented by deliberate judgment-proofing and potential solutions, see Shavell, supra note 5, at 230–32.

[77] We assume for purposes of our consideration of the judgment-proof problem that the firm, presumably insolvent and in bankruptcy proceedings, lacks assets to pay for liability imposed under SOS even if the agency's claim is given priority over those of secured creditors. Although we describe here a number of factors that persuade us that judgment-proof costs are unlikely to be a significant problem under SOS, we acknowledge that the empirical questions underlying these considerations remain open. We are unaware of any studies or data sets that would permit closer empirical analysis of the question at this time.

[78] See Shavell, supra note 5, at 232 (explaining that one benefit of a Pigouvian tax is that it is less susceptible to judgment-proof problems because it determines liability on the basis of risk instead of actualized harm).

[79] Because SOS calls for the application of a single liability determination to all of the firm's regulated sources, firms with fewer sources will be exposed to less variability in liability outcomes and thus a lower probability that the liability outcome will exceed the firm's assets. For example, to determine liability for a firm with just two sources, regulators multiply the liability determination at the selected source by a factor of only two. For a firm with ten sources, however, the liability determination at the selected source is multiplied by a factor of ten to determine the firm's aggregate liability.

[80] For example, an agency applying a periodic inspection policy may divide a large firm's sources into smaller groupings for SOS monitoring. A firm operating 20 sources that have come online in five different time periods might thus face monitoring effectively divided into five separate phases. Consequently, the firm would not be exposed to the liability outcome at any particular source multiplied by 20; rather, the firm would annually be exposed to the liability outcome at each source to be monitored that year multiplied by 4.

ability-insurance regulations.[81] Finally, of course, regulatory liability may not be of sufficient magnitude to raise a realistic possibility that aggregate liability will exceed the firm's assets.[82]

Nevertheless, if a judgment-proof problem arises in a particular case, the agency can address it by providing an option for the agency to sample more than one source. Here, it would be the agency's option, exercisable before any sampling occurs, to determine the number of sources to be sampled.[83] In the event that the agency exercises this option, the average liability determination at sampled sources would be applied perforce as determinative of liability for all of the firm's regulated sources. The firm would be required to pay for the costs of additional sampling.[84]

In the same way that wider sampling reduces risk-bearing costs for firms, sampling additional sources reduces variability in liability outcomes and therefore the likelihood that an outlier will exceed the firm's assets, rendering the firm judgment-proof. Consider

[81] See, e.g., Shavell, supra note 5, at 232; see also 6 U.S.C. §§ 443(a)(1)–(2) (Supp. IV 2005) (requiring firms that produce designated anti-terrorism technologies to carry a reasonable amount of liability insurance).

[82] For example, for violations of the Clean Air Act, the EPA may not impose more than $25,000 per day per violation, and administrative penalties may not exceed $200,000 total. 42 U.S.C. §§ 7413(b)–(d) (2000); Michael Herz, Structures of Environmental Criminal Enforcement, 7 Fordham Envtl. L.J. 679, 711 n.108 (1996) (noting that, in practice, administrative penalties imposed for Clean Air Act violations are relatively small). In fiscal year 2001, for example, the EPA imposed 1,584 administrative penalties totaling $24 million, see EPA, Fiscal Year 2001 Enforcement and Compliance Assurance Accomplishment Report 66, 72, http://www.epa.gov/compliance/resources/reports/planning/results/fy01accomplishment.pdf, an average penalty of just over $15,000.

[83] Where both the firm and the agency elect wider sampling—the agency to avoid the judgment-proof problem and the firm to address risk-bearing costs—as a functional matter the agency will select the minimum number of sources to be sampled, subject to the firm's option for still wider sampling. For example, where the agency has exercised its option to sample three sources, the firm could elect even wider sampling of four sources, but would not be permitted to choose sampling of just two sources. This ensures that the agency will be able adequately to address judgment-proof problems by opting for wider sampling.

[84] As we have noted, this is consistent both with optimal deterrence theory and with the existing practices of the EPA, which requires firms to pay for any reasonable monitoring costs. See supra note 45. The obligation to pay for wider sampling required by the agency should induce firms to be forthcoming with data regarding their ability to pay for potential liability outcomes under SOS. In any event, the agency's option for wider sampling will not ordinarily impose additional information costs on firms, as they must routinely disclose their assets and potential regulatory liability in securities filings or other disclosures required to obtain financing.

again the firm in the previous example with three sources, *A*, *B*, and *C*, producing risk of 10, 20, and 30 units, respectively, and assume that the firm is strictly liable for $1 for each unit of risk. Assume further that the firm has $80 in assets. Under IM, aggregate expected liability is $60; because the firm has adequate assets to pay for its total liability, its incentives for compliance are undiluted. Under SOS, the firm anticipates liability of $30, $60, or $90 when source *A*, *B*, or *C* is randomly selected, respectively; because it anticipates each outcome with 33% probability, its aggregate expected liability remains $60. But if source *C* is selected, the firm will be unable to pay the $90 fine and will pay only $80. Thus, under SOS, the firm internalizes aggregate expected liability of $56.67,[85] and its incentives for compliance may be diluted by the possibility of being judgment-proof.

Now assume that the agency exercises its option to select more than one source for random sampling. Based upon the variance of the firm's liability outcomes and firm wealth, the agency could recognize that the firm's expected liability outcomes may reach $90, exceeding the firm's assets of $80, and require sampling of at least two sources.

As explained in Section III.A above, when two sources are sampled, the firm faces liability outcomes of $45, $60, and $75.[86] By increasing the number of sources to be randomly sampled to two, the agency ensures that no outcome exceeds the $80 maximum the firm is able to pay.[87] Thus, sampling two sources in this case eliminates the judgment-proof problem. By giving regulators the option to require additional sampling to mitigate the judgment-proof problem, agencies can implement SOS without any risk of diluting deterrence.

CONCLUSION

In this Essay, we advance a novel method of administrative enforcement. Our approach permits agencies to duplicate the deter-

[85] This is simply the sum of (33% x $30), (33% x $60), and (33% x $80).

[86] See supra note 72.

[87] Of course, if the firm instead had total assets of only $70, the agency could require sampling of all three sources in order to eliminate the possibility that $75 in liability will be imposed.

rence achieved by cumbersome, expensive iterative monitoring at a fraction of the cost. We have also shown that SOS applies comprehensively, whether a regulation imposes liability on the basis of strict liability or negligence; whether the efficient mode of compliance is standardized or customized; and whether a standard of care is set for each source separately or for a group of sources in the aggregate. We have also shown that while in principle problems of risk-bearing and judgment-proof costs may arise under SOS, in reality these problems are unlikely and in any event can be avoided by minor modifications to SOS that allow for wider sampling with the additional cost charged to the firm.

Beyond indicating the potential benefits of SOS for the EPA, we have described the merits of our proposal only in general terms. We note, however, that there is a wide range of regulatory contexts in which SOS may prove an effective replacement for iterative monitoring. Recently, for example, the Department of the Interior reported that, despite evidence of significant underpayments of oil and gas royalties by lessees of federal lands, the number of audits of such payments fell 22% between 2000 and 2005, due in part to a lack of funding for necessary enforcement personnel.[88] Indeed, the Department estimates that auditing all of the royalties due on those leases would require a tenfold increase in funding.[89] Using SOS, the Department could significantly improve deterrence without requiring any additional funding.

Similarly, the FDA, charged with inspecting drug manufacturing plants throughout the United States, is required by law to inspect such plants at least once every two years.[90] Due to resource constraints, however, the agency has conceded that this level of enforcement is rarely achieved,[91] leading to congressional calls for ad-

[88] Dep't of the Interior, Audit Report of the Minerals Management Service's Compliance Review Process 32–34 (2006), http://www.doioig.gov/upload/2007-G-00011.pdf; Edmund L. Andrews, Report Says Oil Royalties Go Unpaid, N.Y. Times, Dec. 7, 2006, at C1.

[89] See Dep't of the Interior, supra note 88, at 34.

[90] 21 U.S.C. § 360(h) (Supp. IV 2005); 21 C.F.R. § 600.21 (2007).

[91] See, e.g., Alison Young, Eye-Care Plant Uninspected Since '03, Atlanta J-Const., Apr. 19, 2006, at A1 ("Because it lacks resources and inspectors, the FDA manages to inspect plants only about once every five or six years, sometimes even less often, according to agency officials and budget documents.").

ditional enforcement resources to improve deterrence.[92] SOS could permit the FDA to meet its monitoring schedule within current budget constraints.

SOS could also be applied to the civil liability context, reducing litigation costs by permitting courts to determine liability in a single case and apply the results perforce to other cases pending against the same firm. Consider, for example, application of SOS to the thousands of civil liability suits against the U.S. Postal Service arising from traffic accidents.[93] Another application of SOS would promote the use of class actions by permitting courts to determine aggregate liability and damages by resolving one randomly selected claim and applying the outcome perforce to all class members' claims. This approach would avoid costly adjudication of choice of law and other claim-specific legal and factual questions that courts decry as an overwhelming burden and that often serve as a basis for denying otherwise appropriate class certification.[94]

Beyond its practical value in increasing the efficiency of administrative enforcement, SOS serves as a critical heuristic for evaluating the use of individualized decisionmaking in regulation. The lesson of deterrence theory, as illustrated by SOS, is that determining aggregate liability once-and-for-all is just as effective a means of controlling risktaking as determining liability piecemeal. The sole difference is that the former approach is far less expensive. But the savings are not merely financial. Individualized regulation, SOS shows, may unnecessarily squander law enforcement resources, compromising deterrence and consequently jeopardizing the security and well-being of all in society.

---

[92] See, e.g., Minority Staff of H. Comm. on Gov't Reform, 109th Cong., Prescription for Harm: The Decline in FDA Enforcement Activity (2006), http://oversight.house.gov/Documents/20060627101434-98349.pdf.

[93] See Federal Tort Claims Act ("FTCA"), 28 U.S.C. § 1346 (2000); see also Irwin M. Gottlieb, A New Approach to the Handling of Tort Claims Against the Sovereign 1–3 (1967) (noting that the FTCA itself was enacted as a response to the mounting number of tort claims against the government).

[94] See, e.g., In re Bridgestone/Firestone, Inc., 288 F.3d 1012, 1018–19 (7th Cir. 2002). For a proposal to use an SOS-type model for class actions, see David Rosenberg, Overcoming the Choice of Law Barrier to Multi-State Class Actions: A New Method of Sampling (unpublished manuscript, on file with the Virginia Law Review Association).