Published in final edited form as:

Am J Health Econ. 2016 ; 2(1): 66–95. doi:10.1162/ajhe_a_00034.

RISK CORRIDORS AND REINSURANCE IN HEALTH INSURANCE MARKETPLACES:

Insurance for Insurers

TIMOTHY J. LAYTON,

Department of Health Care Policy, Harvard Medical School

THOMAS G. MCGUIRE, and

Department of Health Care Policy, Harvard Medical School

ANNA D. SINAIKO

Department of Health Policy and Management, Harvard School of Public Health

Abstract

Health Insurance Marketplaces established by the Affordable Care Act implement reinsurance and risk corridors. Reinsurance limits insurer costs associated with specific individuals, while risk corridors protect against aggregate losses. Both tighten the insurer's distribution of expected costs. This paper compares the economic costs and consequences of reinsurance and risk corridors. We simulate the insurer's cost distribution under reinsurance and risk corridors using data for a group of individuals likely to enroll in Marketplace plans from the Medical Expenditure Panel Survey. We compare reinsurance and risk corridors in terms of risk reduction and incentives for cost containment. We find that reinsurance and one-sided risk corridors achieve comparable levels of risk reduction for a given level of incentives. We also find that the policies being implemented in the Marketplaces (a mix of reinsurance and two-sided risk corridor policies) substantially limit insurer risk but perform similarly to a simpler stand-alone reinsurance policy.

Keywords

health insurance; risk aversion; adverse selection

I. Introduction

It is well known that health insurance markets are subject to a variety of distortions and market failures that threaten efficiency. The most commonly cited problems are adverse selection and moral hazard. However, an additional factor, the level of risk borne by insurers (i.e., the probability of randomly getting a high-cost draw of enrollees) can also limit efficiency. Under high levels of uncertainty, risk-averse insurers will charge higher premiums, potentially exacerbating inefficiencies from adverse selection (Einav and Finkelstein 2011). Risk-averse insurers may also be less likely to enter high-risk markets, worsening inefficiencies from imperfect competition. It may, therefore, be socially

beneficial to limit insurer risk by implementing reinsurance and risk corridors. While these policies protect insurers from uncertainty, they also weaken insurer incentives to restrain enrollee spending, a trade-off parallel to the trade-off between risk protection and moral hazard described in Zeckhauser (1970) for individual health insurance. Despite widespread use of these policies to limit insurer risk in the new Health Insurance Marketplaces created under the Patient Protection and Affordable Care Act (ACA), Medicare Part D, and many state Medicaid Managed Care markets, there is surprisingly little research on the implications of this trade-off for insurance markets.

Four primary tools protect health plans against the financial risk associated with adverse selection: pricing of health plan premiums, risk adjustment, reinsurance, and risk corridors.¹ All four of these tools are currently being used in the Health Insurance Marketplaces. Each policy ameliorates selection problems differently. Two of these tools deal with “systematic” or predictable risk, typically known in health-care policy as “health risk.” Pricing of health plan premiums can be adjusted for enrollee age or health status so as to match plan revenues for an individual more closely to expected costs. Risk adjustment redistributes revenues from plans with predictably healthier than average enrollees to plans with predictably sicker than average enrollees. The other two tools deal with risk commonly discussed in the actuarial literature, often termed “nonsystematic,” “random,” or “aggregate pricing risk.” This type of risk is generated by the inherent uncertainty around who will enroll in a plan and what health events those enrollees will have during the year. Reinsurance reduces this type of risk by pooling the costs of the sickest enrollees across plans. Risk corridors redistribute revenues from plans earning large profits to plans incurring large losses.²

While there has been a great deal of research on systematic risk and the consequences of risk adjustment and pricing of health plan premiums on insurer and consumer behavior (see McGuire et al. 2013; Glazer, McGuire, and Shi 2013; Handel, Hendel, and Whinston 2013; Shi 2013; Layton 2014; Mahoney and Weyl 2014), nonsystematic risk and the consequences of reinsurance and risk corridors have largely been overlooked. Because nonsystematic risk is likely to be more severe in new markets where insurers have limited experience generating increased uncertainty around the costs of their enrollees, the Health Insurance Marketplaces include the protection of reinsurance and risk corridors only during their first three years. Under current law, after 2016 there will be no protection against nonsystematic risk faced by insurers.

In this paper, we study the effects of the policies designed to limit the nonsystematic risk faced by insurers on insurer risk and incentives. We start by developing a framework for evaluating the performance of reinsurance and risk corridors according to their effect on insurer risk and insurer incentives to restrain enrollee spending. The motivation for limiting insurer risk comes from an assumption that insurer behavior reflects the risk aversion of the managers making decisions at the firm.³ This leads to higher prices and limited competition

¹One could add the individual mandate to this list, a requirement designed to ensure both low- and high-risk individuals participate in the risk pool.

²In the insurance industry, sometimes the terms reinsurance and risk corridors are used interchangeably. We define reinsurance as a policy for which plan reimbursement to an insurer is based on *individual-level* spending, and we define risk corridors as a policy for which plan reimbursement is based on *plan-level* spending.

in the presence of uncertainty. However, while it may be socially desirable to limit insurer risk, doing so may conflict with the goal of encouraging insurers to limit their enrollees' consumption of medical services.⁴ Similar to individual health insurance policies, policies that protect insurers from risk may reduce the insurer's marginal price of an extra dollar of an enrollee's medical spending, weakening insurer incentives to limit enrollees' utilization of medical services (Zeckhauser 1970). Thus, as with individual insurance policies, when designing policies to protect insurers from risk, one must trade off the benefits of insurer risk reduction against the cost of distorting insurer incentives to limit enrollee spending.

We use our framework to compare the relative performance of the two policies used to limit insurer nonsystematic risk: reinsurance and risk corridors. Reinsurance and risk corridors attack the same problem with two fundamentally different strategies. Reinsurance protects insurers against the possibility of randomly enrolling an extraordinarily high-cost *individual*, at the cost of eliminating the insurer's incentive to limit the spending of that individual. Risk corridors, on the other hand, protect insurers against the possibility of randomly enrolling an extraordinarily high-cost *population*, at the cost of eliminating the insurer's incentive to restrain spending on all enrollees *if* it enrolls that high-cost population. The question of which policy provides more risk protection with the strongest incentive to limit enrollee spending is thus an empirical one, depending on the insurer's cost distribution and the distribution of individual medical spending in the population.

We study this question empirically through simulation. We first operationalize the concepts of insurer risk and the insurer's incentive to limit spending among its enrollees by simulating the distributions of potential costs faced by an insurer. We generate the simulated cost distributions by taking a large number of draws from a sample of Marketplace-eligible individuals constructed from the Medical Expenditure Panel Survey (MEPS). Each draw represents a potential state of the world faced by the insurer. We then use the cost distribution to simulate insurer costs, revenues, and ultimately profits under various plan payment systems incorporating reinsurance and risk corridors. Simulations generate measures of insurer risk common in economics and finance. The simulations also provide us with all of the necessary components to evaluate insurer incentives to restrain spending among their enrollees. We measure this incentive empirically using the concept of the "power" of a payment system, defined as the average portion of the marginal dollar of the medical spending of an insurer's enrollees borne by the insurer, as applied by Geruso and McGuire (2014). Thus, if a contract is high powered, insurers have a strong incentive to constrain utilization, and if a contract is low powered that incentive is weak. We measure power empirically via our payment system simulations, and construct this measure by calculating the average change in plan revenues for a given change in plan costs for each of the payment system simulations.

We compare reinsurance and two types of risk corridors, one type where plans that enroll an unexpectedly high-cost group get positive transfers and plans that enroll an unexpectedly

³We provide justification for this assumption in Section IIIA.

⁴This is because when a consumer's marginal price for medical services is reduced by insurance, insured consumers "over-consume" services. This behavior is known as moral hazard, and it is socially desirable for insurers to limit it.

low-cost group are assessed penalties (two-sided risk corridors), and another type where there are no penalties (one-sided risk corridors).⁵ In all simulations, we fix power across policies and observe the level of risk the insurer faces under each policy. We find that for a given level of power, the level of risk reduction caused by a simple reinsurance policy is comparable to the levels of risk reduction generated by simple one-sided or two-sided risk corridor policies. Additionally, we find that a stand-alone “full coverage after a deductible” reinsurance policy generates risk reductions comparable to the set of policies implemented in the Health Insurance Marketplaces during 2014.

The result that reinsurance does as well as risk corridors at reducing insurer risk may seem at odds with economic theory: risk corridors eliminate the tails of the insurer’s cost distribution and, as Arrow (1963) pointed out, an insurance policy that achieves the greatest risk reduction for a given actuarial value would do just that. Reinsurance, on the other hand, results in payments transferred to insurers in all parts of the cost distribution: whether the insurer incurs devastating losses or makes extraordinary profits, if its enrollment base includes at least one high-cost individual it will receive a payment. It seems odd, then, that an insurance policy that pays out when the insurer makes profits (reinsurance) could match (and sometimes surpass) the performance of a policy that only pays out when the insurer incurs large losses (risk corridors). The intuition behind this result is similar to the intuition set out by Zeckhauser (1970): efficiency depends not only on risk protection but also on any behavioral response of the insured, that is, moral hazard. A significant portion of insurer nonsystematic risk is driven by a few extraordinarily high-cost individuals. Because reinsurance focuses on these high-cost individuals, it generates dramatic reductions in insurer risk while only affecting insurer incentives for an extremely small part of its enrollment base. This allows reinsurance to achieve significant risk reduction with only small effects on power. Comparing reinsurance and risk corridors for a given level of power shows that, despite the ability of risk corridors to precisely target the parts of the insurer’s cost distribution that are most influential on insurer risk, reinsurance can perform just as effectively because of its more limited impact on power.

Our results suggest that if policy makers desire to continue to protect insurers in the Marketplaces from nonsystematic risk, they can do so with a simple stand-alone reinsurance policy rather than using more controversial risk corridors. This is an important finding because, as we discuss in Section VI, reinsurance may provide additional benefits such as limiting insurer incentives to engage in inefficient selection, while risk corridors may provide additional distortions, such as encouraging “invest-then-harvest” strategies (Ericson 2014) and exploiting the ability of integrated insurer-providers to manipulate “transfer prices” between providers and insurers.

The paper proceeds as follows. In Section II we explain the various risk-reducing policies being implemented in the ACA Health Insurance Marketplaces. In Section III we set forth a conceptual framework outlining an optimal insurer risk-reducing policy, including consideration of insurer incentives to restrain enrollee medical spending. Section IV

⁵While with two-sided risk corridors, the transfers to unlucky plans are financed by the transfers away from the lucky plans, with one-sided risk corridors the transfers are financed by an actuarially fair per capita risk corridor premium assessed to all plans.

discusses the simulations we use to (1) compare the performance of reinsurance and risk corridors and (2) assess the ability of the proposed reinsurance and risk corridor policies to reduce insurer risk. We also describe the data set of individuals likely to enroll in a Marketplace plan that we construct from the MEPS data set. In Section V we present the results of our simulations. In Section VI we discuss the results, their implications, and their limitations.

II. Policy Background

A. REINSURANCE

We define reinsurance as any policy that reimburses a payer for individual-level spending beyond a preset threshold. Private reinsurance has been available in the US health insurance market for many years, and government-sponsored reinsurance has been implemented at the national level (e.g., the Medicare Part D program) and the state level (e.g., New York, Idaho). Typically, private reinsurance policies cover only the highest-cost cases, while government-sponsored reinsurance reimburses costs starting at much lower thresholds (Bovbjerg et al. 2008; Swartz 2006). In Part D for example, Medicare reimburses health plans for 80 percent of any spending on prescription drugs above an enrollee's catastrophic threshold (MedPAC 2012).⁶ In 2012 Medicare's reinsurance payments to plans amounted to \$14.8 billion, 24 percent of total Medicare payments to prescription drug plans.

Section 1341 of the ACA creates a reinsurance program for the first three years of the Marketplaces, from 2014 to 2016. For 2014, this program reimburses 80 percent of a health plan's annual costs for an enrollee above an "attachment point" of \$60,000 and up to a \$250,000 cap (HHS 2012). Plans are expected to have commercial reinsurance covering costs above \$250,000. All covered claims, not just claims for the federally determined essential health benefits, are eligible for reinsurance (Winkleman et al. 2012). This program is financed by a small reinsurance premium, determined by the Treasury and set annually, that is assessed for all covered lives in non-grandfathered health plans in the United States, including some self-funded plans.

Previous empirical research relies on simulations (as opposed to evaluations of existing programs) to study how reinsurance reduces a plan's potential losses from enrolling high-risk individuals. In a State Children's Health Insurance Program (SCHIP)-eligible population, Sappington et al. (2006) simulate plan profits under varying reinsurance parameters, and find that public reinsurance that fully reimburses a plan for any individual-level expenditures exceeding \$10,000, much lower than the attachment point being implemented in the Marketplaces, reduced average plan losses by 40 percent. An additional line of research analyzes the effects of reinsurance on individual-level, rather than plan-level, costs and profits. Dow, Fulton, and Baicker (2010) use data from a Medicare population to demonstrate that even with reinsurance policies reimbursing insurers for spending incurred by individuals in the top decile of spending, insurers could still expect to lose \$5,400 per individual in the highest risk group, and \$1,700 per individual in the second-highest risk group. Thus, incentives for insurers to avoid enrolling high-cost patients remain.

⁶In 2014, the catastrophic threshold was \$4,550.

Zhu et al. (2013) show that a reinsurance policy that mimics the ACA Marketplace policy of reimbursing insurers for individual-level annual spending above \$60,000 substantially improves the “fit” of payments to costs at an individual level. Finally, Geruso and McGuire (2014) evaluate the performance of the same reinsurance policy as Zhu et al. (2013) with respect to the fit, power, and balance of the payment system. Fit is defined as the R^2 from a regression of each individual’s costs on the revenue that a plan receives for enrolling her. Power is defined as one minus the portion of the marginal dollar of care reimbursed by the regulator. Balance is defined as the variance of service-specific measures of power. They find that reinsurance achieves higher “fit” than risk adjustment for any given level of power.

B. RISK CORRIDORS

We define risk corridors as any policy that reimburses a payer for *plan*-level spending beyond a preset threshold, with the threshold typically being a percentage of the premium charged by the insurer. Risk corridors, also known as aggregate stop-loss reinsurance, have also existed for quite some time in both the private and public sectors (Bovbjerg et al. 2008; Swartz 2006). In the private sector, aggregate (i.e., risk corridor) protection is included in some reinsurance contracts (Bovbjerg et al. 2008). In the public sector, Arizona’s “Health-care Group” program began in the mid-1980s and included a risk corridor-like policy that reimbursed health plans for costs exceeding an aggregate medical loss ratio of 86 percent annually (AcademyHealth 2007). Symmetric risk corridors that limit plans’ profits or losses are also included in Medicare payments to prescription drug plans in Part D, where, after risk adjustment and reinsurance payments for expenditures by high-cost individuals, a plan’s losses or profits can trigger risk corridor payments or penalties (MedPAC 2012). Payments to Part D plans for greater than expected costs are financed by recouping funds from plans with greater than expected profits, and the expected net cost of the program to Medicare is zero.

Section 1342 of the ACA establishes a symmetric risk corridor program for the Marketplaces to operate from 2014 to 2016. The program mimics the Part D policy. Under this program a “target amount” of medical expenditures will be calculated for each health insurer’s covered risk pool, equivalent to their total premiums collected minus an allowed amount for administrative costs and profits. If a plan’s actual expenditures for medical care for its enrollees exceed the target by at least 3 percent, the plan will receive a payment from the risk corridor program. This program is “symmetric” because if a plan’s actual medical expenditures are lower than the target by 3 percent or greater, the plan must make a payment to the risk corridor program.

The Medicare Shared Savings and Pioneer Accountable Care Organization (ACO) Programs also include payment models with risk corridors, allowing ACOs to choose between a one-sided and two-sided arrangement. Under the one-sided arrangement, ACOs share up to 50 percent of savings in excess of a minimum loss ratio (Boyarsky and Parke 2012). Under the two-sided model, ACOs share in 60 percent of savings, but they will also be liable for up to 60 percent of costs above expectation. The vast majority (98 percent) of ACOs have chosen the one-sided model (Centers for Medicare and Medicaid Services 2013).

III. Insurance Principles

A. RISK AVERSION AND RISK NEUTRALITY

In conventional theory of the firm, the firm is assumed to be risk neutral, the rationale being that investors are able to protect themselves against any firm-specific risk by diversification in their investment portfolio.⁷ In practice, however, while firms are owned by shareholders who can diversify to minimize their risk, they are managed by employees. When managers are risk-averse and their incomes are tied to the value of the firm, firms may behave as if they are risk-averse. As a result, the optimal manager contract involves some portion of pay being tied to firm performance, with the exact portion being determined by the trade-off between motivating managers to maximize the value of the firm and minimizing changes to firm behavior that stem from manager risk aversion (Fama and Jensen 1983). Empirical evidence on the correlation between manager pay and firm value suggests that these types of contracts have increased in prevalence over time, largely because of the increasing importance of stock options (Hall and Liebman 1998). Correlation between manager pay and firm value implies that managers may seek to avoid risk.⁸

Risk assumes special significance for firms in the insurance industry. Financial risk from health-care expenses creates demand for health insurance. The law of large numbers reduces the risk of large average losses at the insurer level by pooling the risks of many risk-averse individuals, improving welfare while simultaneously decreasing the insurer's own uncertainty with each additional enrollee. However, some risks are not mitigated by large numbers. For example, in new markets, such as the ACA Marketplaces, the properties of the risk pool may not be well known. Under these circumstances managers at health insurance companies also face the risk of mispricing insurance policies or selecting an unexpectedly high-cost group of initial enrollees.⁹ Following from the discussion above, while owners can hedge away risk, firm managers are likely to be making decisions regarding the pricing of contracts and whether to enter new markets. Because managers' pay is tied to firm value (Hall and Liebman 1998), the firm's choices are likely to reflect manager risk aversion.

In practice, insurers exhibit risk-averse behavior in other ways as well. Risk corridor-like insurance products have historically been purchased by insurers in the private market (Bovbjerg et al. 2008; Swartz 2006). There would be no market for such policies if insurers were risk neutral. The loading factor (the difference between medical claims paid by the insurer and premium revenue) has been estimated to be around 40 percent higher for very small groups than for very large groups (Gruber 1998). While this difference could be due to large fixed costs of insuring a group of individuals, at least some portion of it may also reflect aversion to the increased uncertainty associated with small groups. Finally, a long line of research in the actuarial literature calculates a "risk premium" in order to incorporate it into the loading factor on an insurance policy (see Kahane (1979) and Christofides and

⁷See Rothschild and Stiglitz (1976) for justification of this assumption in the context of insurance. Bulow (2004) presents an argument justifying this assumption in the context of the risk of pharmaceutical patent litigation.

⁸Conceivably, managers could hedge against the risk present in their pay packages on their own. However, given that changes in firm value are quite large (Hall and Liebman (1998) report a standard deviation of changes in firm value of around 32 percent or about \$700 million), it seems likely that liquidity constraints will prevent them from fully hedging against this kind of risk.

⁹As has become obvious in recent months, insurers also face political risk where policies and regulations that affect the risk pool can change after insurers set prices.

Smith (2001) for some examples of this literature). Most of these studies relate the “cost” of risk to be priced into the insurance premium to the standard deviation of the cost distribution, a measure we will use later. For all of these reasons, the following analysis takes as a point of departure that the firms considering participating in Health Insurance Marketplaces act as though they are risk-averse.

B. RISK MEASURES

We measure insurer nonsystematic risk in two ways. The first is the standard deviation of the distribution of potential outcomes. This measure is proportional to what is known as the “risk premium,” or the dollar amount an agent is willing to pay to eliminate uncertainty, for an individual with constant absolute risk aversion utility.¹⁰ The second measure of risk is the likelihood of a large loss. This measure is proportional to the risk premium of an individual who faces a large penalty for incurring a large loss. The measure is motivated by managers or owners/investors facing a large penalty for being forced to declare bankruptcy or for depleting the insurer’s reserves.

We define potential outcomes for insurers using the distribution of expected costs of the insurer’s enrollees.¹¹ Our construction of costs includes only medical and pharmaceutical expenditures and ignores administrative costs, as these fixed and predictable costs are nonrandom and do not affect insurer risk. We treat the mean of the cost distribution as the insurer’s expected cost per enrollee, and we assume that any outcome below this value results in insurer “profits” and any outcome above this value results in insurer “losses.”¹² Figure 1 illustrates this cost distribution. Essentially, each year an insurer receives a draw from this distribution. If an insurer receives a draw from the left side of the distribution, it will earn unexpected profits in that year. If the insurer receives a draw from the right side of the distribution, it will experience unexpected losses.

C. THEORY OF INSURANCE FOR INSURERS

The seminal papers establishing the theory of insurance in health economics are Arrow (1963) and Zeckhauser (1970). Arrow establishes that, with a limited budget, the optimal insurance policy is full coverage after a deductible. In Arrow’s model, individuals have diminishing marginal utility in nonmedical consumption and uncertainty about their future health status and, thus, their medical expenses. Under these assumptions, individuals value an extra dollar of coverage against medical expenses more as the potential expense gets larger. Therefore, coverage for large expenses provides larger welfare gains than coverage for small expenses.

In the context of protecting health insurers from risk, if we assume no moral hazard, Arrow’s (1963) principle of the optimality of full coverage after a deductible implies that it is inefficient to use reinsurance to provide risk protection. The logic for this argument is as

¹⁰The risk premium is proportional to the variance of the random variable in the case where the random variable is distributed normally.

¹¹The insurer’s cost distribution is the distribution of the average costs of the insurer’s enrollees. Because the cost distribution is a distribution of averages, the central limit theorem implies that this distribution will be approximately normal.

¹²This definition of “profits” and “losses” comes from an assumption that insurers price their insurance policies at the average expected cost of potential enrollees. Such pricing is implied by an assumption of perfect competition.

follows: because reinsurance protects against individual-level losses, some insurers that *ex post* have an outcome in the left-hand side of the cost distribution (i.e., positive profits) will receive reinsurance payments if they have one or two high-cost cases and many low-cost cases. At the same time, some insurers that *ex post* have an outcome in the right-hand side of the cost distribution (i.e., aggregate losses) will receive no reinsurance payments if they have many cases whose realized costs are only slightly above the expected cost and no high-cost cases. A more efficient policy would mimic Arrow's (1963) optimal insurance policy of full coverage after a deductible by reimbursing only plans incurring large aggregate losses. This type of policy has the benefit of using fewer dollars to achieve greater reductions in the variance of the profit distribution and the probability of a large loss for plans. This is exactly how risk corridors operate.¹³

D. POWER AND INCENTIVES FOR EFFICIENCY

Insurer risk reduction, however, is clearly not the only, or even the principal, objective of the regulator. The regulator may also wish to encourage insurers to use the tools available to them to restrain spending in order to keep prices low. It is well known that cost-based provider payment systems, also known as fee-for-service systems, are likely to result in overuse of medical services (Ellis and McGuire 1986). Cost-based insurance market payment systems are likely to have similar undesirable effects. Most policies that reduce insurer risk, including both reinsurance and risk corridors, result in payment systems that are partially cost-based. Thus, the regulator faces a trade-off between reducing insurer risk and providing insurers with the incentive to control costs.

To incorporate this trade-off into our framework, we draw on contract theory and use the concept of the *power* of a payment system (Geruso and McGuire 2014; Laffont and Tirole 1993). Under a contract, when a principal's payments to an agent are tied to the agent's costs so that the insurer's marginal price of an additional dollar of an enrollee's medical spending is reduced, the contract is low powered, weakening an insurer's incentives to control costs. When payments are independent of costs, however, the contract is high powered, with stronger incentives for cost control.

We follow Geruso and McGuire (2014) and define power as the share of costs borne by a health plan at the margin. For an individual enrollee, this share is characterized by the derivative of the payment that a plan receives for enrolling the individual with respect to her utilization, $\frac{dp_i}{dx_i}$. In a setting where an insurer is fully reimbursed for each dollar spent on an enrollee's care, this derivative will be one. In a setting where a plan is paid a fixed capitation fee for each enrollee, on the other hand, the derivative will always be equal to zero. At a population level, we define power (ρ) as

¹³Risk adjustment is another policy related to insurer risk that is often discussed along with risk corridors and reinsurance. However, while risk adjustment has important effects on the mean of an insurer's distribution of expected costs, its effects on the variance of the distribution are minimal. Because all of our measures of risk are based in some way on the variance of the cost distribution and are unrelated to the mean of the distribution, we abstract from risk adjustment through much of this paper.

$$Power \equiv \rho = 1 - \frac{1}{N} \sum_{i=1}^N \frac{dp_i}{dx_i}. \quad (1)$$

This measure of power is thus the share of the marginal dollar of an individual's spending borne by the insurer averaged over all individuals in a plan. A fully high-powered contract has $\rho = 1$, whereas a completely low-powered contract has $\rho = 0$. Here we assume that more power is better and that any sacrifice of power away from 1.0 introduces inefficiency. While this assumption may not always be true in practice, it provides a reasonable benchmark for analyzing the relative efficiency of risk-reducing policies.¹⁴

Both reinsurance and risk corridors tie an insurer's payments to its costs (at least in part), and thus affect power. While risk corridors follow Arrow's (1963) prescription to eliminate the parts of the insurer's cost distribution that contribute the most to insurer risk, they do so by setting power to zero *for every enrollee* in any state of the world where the insurer's costs exceed the risk corridor "deductible." Reinsurance, on the other hand, is not as well targeted to the tails of the cost distribution that are most important for insurer risk, compensating insurers whose draws from the distribution are close to the mean as well as those whose draws are in the tails of the distribution. However, it does focus on the high-cost individuals that contribute the most to insurer risk, and it only affects power for the small number of enrollees that incur extremely high costs. In theory, the optimal policy would likely compensate insurers for enrolling extremely high-cost individuals only in states of the world where the insurer's draw from the distribution results in large losses. In practice, however, such a policy seems impractical. Thus, it is important to determine whether reinsurance, with its limited effect on both power and insurer risk, can provide equivalent (or better) risk reduction for a given level of power than risk corridors, with their strong effect on both power and risk. This is an empirical question that we investigate in the next sections.

Also of interest is the relative effectiveness of one-sided versus two-sided risk corridors in terms of the risk-power trade-off. In a two-sided risk corridor, risk corridor payments to Marketplace plans incurring large losses are funded through transfers from plans earning significant profits. A one-sided risk corridor is an alternative funding arrangement where plans incurring large losses receive the same payments as they would under a two-sided risk corridor, but the payments are funded by an actuarially fair per capita risk corridor premium. The two types of risk corridors differ in their implications for insurer risk and power. With respect to the variance of the cost distribution, two-sided risk corridors have twice the effect on insurer risk as one-sided risk corridors. However, two-sided risk corridors also have twice the effect on power, so it is unclear whether two-sided or one-sided risk corridors provide more risk reduction for a given level of power. With respect to the probability of a large loss, our other measure of insurer risk, it is clear that one-sided risk corridors will dominate two-sided risk corridors. Both offer the *same* level of risk reduction, but the effect of two-sided risk corridors on power remains twice the size of the effect of one-sided risk

¹⁴While it is clear that a system where insurers are fully reimbursed for every dollar spent on their enrollees' care provides too little incentive to constrain utilization, a system where insurer payments and costs are independent may lead to *under*-utilization (Ellis and McGuire 1986; Newhouse 1996). In this case the optimal level of power will be less than one, and analyses like those conducted here can help identify the risk corridor/reinsurance policies that minimize risk for any desired payment system policy.

corridors. We will also compare the effectiveness of these two policies in our simulations, though intuition suggests that one-sided risk corridors will at least weakly dominate two-sided risk corridors.¹⁵

IV. Data and Methods

In the following empirical analysis of these policies we first simulate a distribution of outcomes for plans participating in a Marketplace to calculate baseline risk measures for the base case with no risk-reducing policies. We then apply reinsurance and risk corridors to the distribution, recalculate the risk measures, and calculate our measure of power under each policy to determine how each policy affects insurer risk and the incentive for insurers to limit spending among their enrollees.

A. DATA ON THE MARKETPLACE POPULATION AND HEALTH-CARE SPENDING

The Medical Expenditure Panel Survey (MEPS) is a large, nationally representative survey of the civilian noninstitutionalized US population with information on approximately 33,000 individuals annually. Summary statistics for the sample can be found in Table 1. We identify a Marketplace-eligible population following methods in McGuire et al. (2013). Pooling MEPS data from Panels 9 (2004–05) through 14 (2009–10), we select a population of individuals and families eligible for enrollment in Marketplaces based on income, insurance, and employment status. Specifically, we select adult, nonelderly individuals (aged 18–64) in households earning at least 138 percent of the federal poverty level (FPL) and children in households with income of at least 205 percent of FPL. Selection criteria into the Marketplace population, as defined by the ACA, include individuals living in households in which an adult was ever uninsured, a holder of a non-group insurance policy, self-employed, employed by a small employer, or paying an out-of-pocket premium for their employer-sponsored health insurance (ESI) plan that is deemed to be unaffordable. If an individual meets the selection criteria in at least one of the two survey years, she is included in the sample. The data set comprises 44,210 Marketplace-eligible individuals, 11,773 of whom have only one year of data and 32,437 of whom have two years of data, generating a total sample of 76,647 person-years.

The MEPS includes data on total health-care expenditures for each individual during each year. It also includes information such as diagnoses from all of the medical events that result in those expenditures (such as office visits, hospital stays, prescriptions filled). MEPS data understate health expenditures (Sing et al. 2002; Aizcorbe et al. 2012; Zuvekas and Olin 2009). Discrepancies are driven by both underreporting of health-care utilization and underrepresentation of high-expenditure cases due to the exclusion of patients who are

¹⁵One might also conceive of a “hybrid” risk corridor policy that funds a portion of payments to plans incurring large losses through a per capita premium and funds a portion through transfers from plans earning large profits. In this paper, we choose to look only at the extreme cases of “pure” one-sided risk corridors and “pure” two-sided risk corridors. We show below that “pure” one-sided and two-sided risk corridors are virtually indistinguishable in terms of the risk-power trade-off with respect to the standard deviation measure of risk. This implies that any hybrid policy would also be indistinguishable. As discussed in the paragraph above, intuition (and the empirical results below) imply that the “pure” one-sided policy will always dominate the “pure” two-sided policy with respect to the value at risk measures of insurer risk (95th and 99th percentiles of cost distribution) because both policies offer the same risk reduction, but the two-sided policy results in lower power than the one-sided policy. This same intuition would imply that any hybrid policy would also be dominated by the “pure” one-sided policy with respect to these measures.

institutionalized or hospitalized longer than 45 days. Zuvekas and Olin (2009) suggest that total expenditures be inflated by a factor of 1.09 for individuals with an inpatient claim and by a factor of 1.546 for all other individuals. We adopt this correction, inflating expenditures of the individuals in our sample as directed.

B. COST DISTRIBUTION

The simulated cost distribution describes an insurer's average cost per enrollee in all potential states of the world and the likelihood of each state. In order to simulate the insurer's cost distribution we first fix the size of the plan at Q enrollees. We take M random samples of Q individuals from our sample of $N = 76,647$ Marketplace-eligible individuals.¹⁶ For each random sample, we calculate the average cost of the Q chosen individuals, $\bar{x}_m$. The set of average costs from all draws from the Marketplace-eligible sample forms our simulated cost distribution,

$$f(\bar{x}) = Pr(\bar{x}) = \frac{\sum_{m=1}^M 1(\bar{x} = \bar{x}_m)}{M}. \quad (2)$$

We use this simulated distribution to test the effects of reinsurance and risk corridors on risk and power.

We acknowledge that using this method to simulate the insurer's cost distribution may neglect certain sources of insurer risk. The insurer's true cost distribution is generated by all types of uncertainty surrounding the insurer's costs. This includes uncertainty about the average mix of enrollees in the market, the portion of the enrollees that choose the particular insurer's plan, changes in medical technology, and the like. By simulating the insurer's cost distribution through resampling from a population of Marketplace-eligible individuals, we are making the implicit assumptions that (1) all of the individuals we identify as Marketplace-eligible are equally likely to enroll in the insurer's plan and that (2) the MEPS claims data accurately characterize each individual's expected cost to the insurer. While these assumptions may not hold in practice, we believe they are reasonable given our data constraints. Additionally, if the assumptions are violated, the conceptual points made in our simulations are likely to hold.

C. REINSURANCE

Our methods for modeling reinsurance mimic the methods used in Zhu et al. (2013). Let x_i be individual i 's total annual cost to the insurer. Let x be the reinsurance threshold above which costs are reimbursed. Finally, let δ be the rate at which costs are reimbursed. The reinsurance payment received by the insurer when enrolling individual i , re_i , is defined as follows:

¹⁶Note that this implies that we treat each person-year as an independent observation. We do this to allow us to simulate the cost distributions of large insurance plans given a limited sample of Marketplace-eligible individuals. Because many individuals are present in the data for two years, each individual's two person-year combinations are not in reality independent observations. Treating them as independent observations will tend to understate the variance of the insurer's cost distribution.

$$re_i = \begin{cases} 0 & \text{if } x_i < \hat{x} \\ \delta(x_i - \hat{x}) & \text{if } x_i \geq \hat{x} \end{cases} \quad (3)$$

We assume that plans' paid claims are equal to total claims; in other words, we assume that the plan covers all health-care costs incurred by an individual during a given year. We also assume that reinsurance is funded through a per capita, actuarially fair reinsurance fee collected for each Marketplace enrollee.¹⁷ The fee, denoted ref , is equal to the average reinsurance payment for the entire population:

$$ref = \frac{1}{N} \sum_i re_i. \quad (4)$$

This ensures that reinsurance will be budget neutral.

We apply reinsurance to the simulated cost distribution as follows. First, we determine re_i for each individual in the Marketplace-eligible sample. We then calculate ref and re-define each individual's cost to the plan as $x_i^{re} = x_i - re_i + ref$. Finally, we simulate the insurer's cost distribution using the same methods described in Section IVB but using the new definition of insurer costs, x_i^{re} . The plan's simulated cost distribution under reinsurance is thus

$$g(\bar{x}^{re}) = Pr(\bar{x}^{re}) = \frac{\sum_{m=1}^M 1(\bar{x}^{re} = \bar{x}_m^{re})}{M}. \quad (5)$$

D. RISK CORRIDORS

Risk corridors transfer funds to an insurer if the sum of covered health-care costs of the insurer's enrollees is greater than a fixed percentage of a target. In the Marketplace risk corridor policy, the target is defined as an insurer's total premium revenues minus some combination of administrative costs and profits. To simplify, we define the target as the mean of the average cost distribution.¹⁸ Thus, if an insurer's realized average per capita cost is greater than the expected average cost (the target) by a defined minimum percentage, the risk corridor will reimburse the insurer for a portion of its costs according to the risk corridor cost-sharing parameters.

To apply risk corridors to the simulated cost distribution, we first construct the cost distribution using the methods described in Section IVB. Let p be the mean of the simulated cost distribution. Thus the target is

¹⁷Note that this is different from how reinsurance will be funded in practice (described in Section II). We do this to allow for an "apples-to-apples" comparison of reinsurance and risk corridors by forcing them both to be "self-funding."

¹⁸Again, only factors that involve uncertainty impact risk corridor payments and our risk measures. Administrative costs and profits affect the level of outcomes but not the other moments of the outcome distribution (i.e., variance), so they don't matter for risk corridor payments or for our risk measures. This assumption could also be justified by assuming perfect competition and no administrative costs.

$$\bar{p} = \frac{1}{M} \sum_{m=1}^M \bar{x}_m. \quad (6)$$

Now we introduce three new parameters: average plan costs as a percentage of the target, $\alpha_m = \frac{\bar{x}_m}{\bar{p}}$; the upper threshold, $\bar{\alpha}$; the lower threshold, $\underline{\alpha}$ and the portion of costs reimbursed, θ . Risk corridor payments are based entirely on the value of α_m implied by the insurer's draw from the expenditure distribution. Specifically, the risk corridor gives the insurer the following transfer:

$$t(\alpha_m) = \begin{cases} \theta(\bar{x}_m - \bar{\alpha}\bar{p}) & \text{if } \bar{\alpha} < \alpha_m \\ 0 & \text{if } \underline{\alpha} < \alpha_m < \bar{\alpha} \\ \theta(\bar{x}_m - \underline{\alpha}\bar{p}) & \text{if } \alpha_m < \underline{\alpha} \end{cases}. \quad (7)$$

Suppose that the risk corridor reimburses plans for 50 percent of costs above 108 percent of the target. In this case, $\bar{\alpha} = 1.08$ and if for a given plan $\alpha > 1.08$ then the plan will receive a transfer of $0.5(\bar{x} - 1.08\bar{p})$. More generally, if an insurer gets a “bad” (“good”) draw from the expenditure distribution such that α_m is large (small), it will receive a positive (negative) transfer from the regulator to offset some portion of its losses.

We calculate the insurer's α_m for each draw m , and use α_m along with the above definition to determine $t(\alpha_m)$. We then redefine the insurer's cost to the plan for enrolling individual i in state of the world m as $x_{im}^{rc} = x_i - t(\alpha_m)$. Finally, we use this new definition of costs to construct the new cost distribution with risk corridors:

$$h(\bar{x}^{rc}) = Pr(\bar{x}^{rc}) = \frac{\sum_{m=1}^M 1(\bar{x}^{rc} = \bar{x}_m^{rc})}{M}. \quad (8)$$

With large M , the cost distribution will be symmetric. If $\underline{\alpha} = 1 - (\bar{\alpha} - 1)$, the risk corridor transfer, $t(\alpha_m)$, will be equal to zero in expectation. In other words, the risk corridor policy is “budget neutral” in expectation.

For the risk corridor policy described thus far, the positive transfers *to* plans with unexpectedly high costs are funded by transfers *from* plans with unexpectedly low costs. As discussed in Section IIID, we refer to this type of risk corridor policy as two-sided risk corridors.

As discussed in Section IIID, the transfers to the plans drawing an unexpectedly high average cost could also be funded via an actuarially fair premium assessed to *all* plans in the market. We refer to this type of policy as a one-sided risk corridor, and we simulate an additional version of the cost distribution that reflects the effects of such a policy. Under a one-sided risk corridor policy, the regulator pays insurers the following transfer:

$$t(\alpha_m) = \begin{cases} \theta(\bar{x} - \bar{\alpha}\bar{p}) & \text{if } \bar{\alpha} < \alpha_m \\ 0 & \text{if } \alpha_m < \bar{\alpha} \end{cases}. \quad (9)$$

Insurers are also assessed an actuarially fair per capita fee, denoted rcf . We redefine α_m to incorporate this fee: $\alpha_m^{os} = \frac{\bar{x} + rcf}{p}$.¹⁹ The fee is equal to the expected value of the transfer, taken over all potential states of the world:

$$rcf = \frac{1}{M} \sum_{m=1}^M t(\alpha_m^{os}). \quad (10)$$

Note that because the transfer is a function of α_m^{rc} , rcf is a function of α_m^{rc} . Recall that α_m^{rc} is also a function of rcf . This implies that for each $\bar{a}$ there is an equilibrium value for rcf that satisfies equation 10.

To apply one-sided risk corridors to the simulated cost distribution, we redefine insurer costs as $x_{im}^{os} = x_i - t(\alpha_m^{os}) + rcf$. Because α_m^{os} is a function of rcf and rcf is a function of α_m^{os} , we use an iterative procedure to find the value of α_m^{os} that causes equation 10 to hold. Finally, we use the new definition of plan costs to construct the new cost distribution under one-sided risk corridors:

$$h^{os}(\bar{x}^{os}) = Pr(\bar{x}^{os}) = \frac{\sum_{m=1}^M 1(\bar{x}^{os} = \bar{x}_m^{os})}{M}. \quad (11)$$

E. ESTIMATION OF RISK

As discussed in Section IIIB above, we quantify insurer risk in two ways. The first is the standard deviation of the expected cost distribution, where risk is increasing in the standard deviation. The second measure characterizes the probability of a large loss. It is termed the “value at risk” in the finance literature, and is defined as the Y th percentile of the expected cost distribution. In our case, we will set $Y = 95$ or $Y = 99$, so larger values of the “value at risk” imply higher risk.

F. ESTIMATIONS OF POWER

As noted above, we measure power as the portion of the marginal dollar spent on an average enrollee’s health care by the health plan:

$$Power \equiv \rho = 1 - \frac{1}{N} \sum_{i=1}^N \frac{dp_i}{dx_i}. \quad (12)$$

For the reinsurance and risk corridor policies we simulate, there is no closed-form solution for this definition of power, so we estimate it via simulation. First, we specify plan revenues, p_{im} , as the sum of plan premiums, net reinsurance payments, and (net) risk corridor payments:

¹⁹The risk corridor policy in the ACA establishes that α be based on insurer costs *net of all other policies*. This implies that α is based on costs after reinsurance is applied or, in this case, after the fee is assessed. Without this condition, the policy would not be budget neutral.

$$p_{im} = \bar{p} + re_i - ref + t(\alpha_m). \quad (13)$$

We assume that premiums are constant for all individuals and set equal to the average cost in the Marketplace-eligible population, as they would be in a perfectly competitive market with zero profits.²⁰

Next, we estimate $\frac{1}{N} \sum_{i=1}^N \frac{dp_i}{dx_i}$ by simulation. The derivative of plan payment with respect to utilization, $\frac{dp_i}{dx_i}$, describes how much plan payments change given a \$1 increase in utilization. Note that this value is defined by the structural rules of the payment system, not by any behavioral parameters. Reinsurance and risk corridor payments are functions of realized costs, so any shift in realized costs will affect these payments. We can thus determine the derivative in each policy environment by simulating a decrease in each individual's spending and then using the payment system rules described above to determine the corresponding change in plan payments. For purposes of the simulation, we assume that when an insurer reduces utilization, the premium, reinsurance fee, and risk corridor target remain constant. We do this because our goal is to characterize the plan's *incentive* to control costs, not to simulate equilibrium outcomes when plans are able to achieve cost reductions. We calculate this derivative for each draw from the cost distribution.

More specifically, for each policy we start by calculating p_{im} for each draw from the Marketplace-eligible sample. We then implement an insurer reduction of 0.1 percent of each individual's health-care costs for each draw, i.e., $x_i^* = 0.999x_i$.²¹ We then re-simulate reinsurance and risk corridors where costs are equal to x_i^* instead of x_i , generating new reinsurance and risk corridor payments (re_i^* and $t(\alpha_m^*)$). Premiums and reinsurance and risk corridor fees remain the same. We use these values to recalculate plan revenues, $p_{im}^* = \bar{p} + re_i^* - ref + t(\alpha_m^*)$. Power is then calculated as

$$\hat{\rho} = 1 - \frac{1}{M} \sum_{m=1}^M \frac{1}{N} \sum_{i=1}^N D_{im} \frac{p_{im} - p_{im}^*}{x_i - x_i^*}, \quad (14)$$

where D_{im} is equal to one if individual i was chosen in draw m and zero otherwise. This characterization captures the expected change in plan revenues for an incremental change in plan costs and leads to a measure of power, the portion of the marginal dollar of an enrollee's health-care spending borne by the plan.

²⁰The ACA allows for premiums to vary by age, geography, and smoking status. We abstract from these allowances here, as they are unrelated to utilization and, thus, to power. Note that in this case the premium is equal to the risk corridor "target."

²¹We use a very small reduction to avoid problems stemming from the fact that a large-to-moderate size reduction has asymmetric effects on the left and right sides of the cost distribution (for risk corridors). This asymmetry results in an overestimation of the power of a payment system, including one-sided risk corridors, and an underestimation of the power of a payment system, including two-sided risk corridors.

G. SIMULATION METHODS

We perform two types of simulations. First, we compare the performance of reinsurance and risk corridors using a simplified form of each policy. For risk corridors, we set $\theta = 1$. Effectively, the simplified policy provides “full coverage after a deductible.” Similarly, we simplify the definition of reinsurance by setting $\delta = 1$. Thus, the policy reimburses plans for 100 percent of an individual’s costs above a threshold. These simplifications compare policies based on a similar reimbursement structure. We allow the reinsurance and risk corridor cutoffs (i.e., $\bar{a}$ and x) to vary over a wide range. For reinsurance, we allow the cutoff to vary between \$0 and \$375,000. For risk corridors we allow the cutoff to vary between 100.1 percent and 130 percent of the target. We also allow risk corridors to be either two-sided or one-sided, and for two-sided risk corridors we set $\underline{a}$ by reflecting $\bar{a}$ across the target: $\underline{a} = 1 - (\bar{a} - 1)$. For each simulation, we calculate our measure of power and each of the risk measures described in Section IVE and IVF. We then compare the level of risk faced by an insurer generated by a reinsurance or risk corridor policy with a given level of power.

The second set of simulations evaluates the performance of the policies implemented in the Marketplaces from 2014 to 2016. We simulate each proposed policy (reinsurance and risk corridors) on its own and combined in the full package of risk-reducing policies. We adapt our general framework to the actual policies being implemented in the Marketplaces. Both the ACA reinsurance and risk corridor policies are slightly more complicated than those modeled above in that they have multiple thresholds and the portion of costs reimbursed varies across thresholds. In the Marketplaces, risk corridor transfers are defined as follows:

$$t(\alpha_m) = \begin{cases} 0.025 p_n + 0.8(c_k - 1.08 p_n) & \text{if } 1.08 < \alpha_m \\ 0.5(c_k - 1.03 p_n) & \text{if } 1.03 < \alpha_m < 1.08 \\ 0 & \text{if } 0.97 < \alpha_m < 1.03 \\ 0.5(c_k - 0.97 p_n) & \text{if } 0.92 < \alpha_m < 0.97 \\ -0.025 p_n + 0.8(c_k - 0.92 p_n) & \text{if } \alpha_m < 0.92 \end{cases} \quad (15)$$

Similarly, reinsurance payments will be paid according to the following definition:

$$re_i = \begin{cases} 0 & \text{if } x_i < 60,000 \\ 0.8(x_i - 60,000) & \text{if } 60,000 < x_i < 250,000 \\ 152,000 + 0.85(x_i - 250,000) & \text{if } 250,000 < x_i \end{cases} \quad (16)$$

All results presented below are from simulations of a plan with 20,000 enrollees, and the cost distribution is simulated with 5,000 draws from the Marketplace-eligible sample. Similar results from simulations of plans with 5,000 and 50,000 enrollees can be found in the Online Appendix (http://www.mitpressjournals.org/doi/suppl/10.1162/AJHE_a_00034).

V. Results

A. REINSURANCE VERSUS RISK CORRIDORS

Figure 2 plots the “risk” of the insurer’s cost distribution for each policy and each risk measure. In all three panels, the y-axis measures insurer risk, and the x-axis shows the

reinsurance (top x-axis) and risk corridor (bottom x-axis) cutoff points.²² The lines plot the measure of insurer risk under each policy for each simulated reinsurance or risk corridor cutoff. Each panel of the figure shows a different risk measure, with the standard deviation on top, the 95th percentile of the cost distribution in the middle, and the 99th percentile of the cost distribution on bottom. For example, the top panel shows that a one-sided (two-sided) risk corridor policy with a cutoff of 1.02, so that plans are reimbursed for all costs above 102 percent of the target, generates a cost distribution with a standard deviation of 53 (46), a 13 percent (25 percent) reduction from the cost distribution in a setting where there is no reinsurance and no risk corridors. This panel also shows that a reinsurance policy with a cutoff of \$100,000 (i.e., plans are reimbursed for any individual-level costs exceeding \$100,000) generates a cost distribution with a standard deviation of 48, a 21 percent reduction from the cost distribution in a setting where there is no reinsurance and no risk corridors.

Figure 2 shows that all of the policies significantly affect insurer risk according to all measures, especially as the reinsurance cutoff approaches \$0 and the risk corridor cutoff approaches one. This is to be expected, as these policies act to transfer risk from the insurer to the government. Also as expected, with respect to the standard deviation measure of insurer risk, two-sided risk corridors provide about twice as much risk reduction as one-sided risk corridors at a given cutoff point, and with respect to the value at risk measures, one-sided and two-sided risk corridors affect risk identically. It is also interesting to note that while risk corridors exhibit a somewhat linear relationship between the cutoff and insurer risk, the marginal return to a lower reinsurance cutoff in terms of insurer risk reduction is somewhat small until it increases dramatically as the reinsurance cutoff approaches zero.

Figure 3 shows the power of each payment system under each policy. Again, the top x-axis shows the reinsurance cutoffs and the bottom x-axis shows the risk corridor cutoffs. The y-axis shows the power of each policy. As expected, the power of both risk corridor policies (reinsurance) approaches zero as the cutoff approaches one (\$0). Additionally, the power of the risk corridor policies (reinsurance) approaches one as the cutoff moves away from one (\$0). Similar to the risk measures, the marginal effect of a lower reinsurance cutoff on power is somewhat small until it increases dramatically as the reinsurance cutoff approaches zero. The power curves for risk corridors also look similar to the risk corridor risk curves from Figure 2, with two-sided risk corridors resulting in a drop in power about twice the size of the drop in power due to one-sided risk corridors for any given cutoff. Comparing Figures 2 and 3 reveals the explicit trade-off between power and insurer risk. In all simulations, as risk is reduced, so is power.

While Figures 2 and 3 provide some indication of the relative effectiveness of each of the policies, they do not allow for apples-to-apples comparisons, especially not for reinsurance versus risk corridors. It is not clear whether a reinsurance policy with a cutoff of \$100,000 is comparable to a one-sided risk corridor policy with a cutoff of 1.10 or 1.03 or 1.35. To allow

²²Note that for risk corridors, the x-axis shows only the top cutoff point. For two-sided risk corridors there is also a bottom cutoff point equal to $\underline{a} = 1 - (\bar{a} - 1)$.

for comparisons, we need to normalize the policies using some consistent value. Moreover, reinsurance and risk corridor policies should be evaluated according to their effects on both risk *and* power. In Figure 4, we combine Figures 2 and 3 to show the risk and power for each simulated policy in the same plot and to allow for comparisons on both metrics. Each panel of the figure presents a different risk measure. In each panel, power is shown on the x-axis and the risk measure is shown on the y-axis. The data in the figure are the same data presented in Figures 2 and 3.

The “best” policies will be the ones closest to the bottom right quadrant of each panel, as these are the policies that generate large risk reductions while keeping power high. With respect to the standard deviation, reinsurance clearly offers the most risk reduction for a given level of power. This is because reinsurance virtually eliminates the effects on insurer risk of the few high-cost individuals who account for a large portion of insurer risk. At the same time, reinsurance has a fairly small effect on power because it only affects the marginal cost to the insurer of an additional dollar of spending for each of the small number of individuals with spending that exceeds the (often quite high) reinsurance cutoff point. This is similar to the reason why reinsurance dramatically outperforms other policies with respect to the “fit” of a payment system (Geruso and McGuire 2014). For the standard deviation of the cost distribution, one- and two-sided risk corridors appear virtually identical with respect to the risk-power trade-off. This implies that the additional risk reduction provided by the two-sided risk corridor is almost exactly offset by additional losses in power.

With respect to the other risk measures, it is less clear which policy is best. For the 95th percentile of the cost distribution, reinsurance dominates at high levels of power while one-sided risk corridors dominate at lower levels of power. For the 99th percentile of the cost distribution, one-sided risk corridors dominate both two-sided risk corridors and reinsurance, though differences across the policies are relatively small, with the maximum difference between the 99th percentile of the cost distribution under reinsurance and one-sided risk corridors at a given level of power being only around \$30 at a level of power of 0.8. It is interesting to note that according to all three risk measures, one-sided risk corridors weakly dominate the two-sided risk corridors used in the Marketplaces from 2014 to 2016, always generating at least as much risk reduction for a given level of power. This result is consistent with the intuition laid out in Section IIID.

B. PROPOSED MARKETPLACE POLICIES

Each of the policies being implemented in the Marketplaces from 2014 to 2016 reduces insurer risk (Table 2). When simulated separately, the Marketplace reinsurance policy reduces the standard deviation of the insurer’s cost distribution by \$17 (28 percent) and the Marketplace (two-sided) risk corridor policy reduces the standard deviation of the insurer’s cost distribution by \$3 (5 percent). When the policies are implemented together, the risk reduction is nearly the same as that with the simple “full coverage after a deductible” reinsurance policy alone. While the Marketplace reinsurance policy has a modest effect on the value at risk, reducing the 95th percentile of the cost distribution by \$28, the Marketplace risk corridor policy has little to no effect on this risk measure.²³ Figure A6 in

the Online Appendix illustrates the effects of these policies on the cost distribution by plotting the kernel densities of the simulated distributions.

The final column of Table 2 reports power under each policy configuration. The Marketplace reinsurance policy results in a much larger reduction in power than the marketplace risk corridor policy. In fact, when the risk corridor policy is combined with the reinsurance policy, power is only slightly lower than with reinsurance alone.

To put the performance of the Marketplace policies in context, Figure 4 plots the measures of insurer risk and power under the full “hybrid” Marketplace policy (reinsurance and two-sided risk corridors) along with the simulated counterfactual policies discussed in Section VA. In each figure, the black square represents the combination of power and risk under the hybrid Marketplace policy. The simple “full coverage after a deductible” reinsurance policy performs comparably to the complicated hybrid Marketplace policy with respect to all three risk measures. This suggests that if policy makers desire to continue to offer insurers the same level of risk protection and the same level of power as the initial “hybrid” Marketplace policy, they can do so with a simple reinsurance policy. The reinsurance policy that produces comparable insurer risk and incentives to limit enrollee medical spending has a cutoff point of \$65,000 and reimburses insurers for all individual-level spending exceeding that value.

VI. Discussion

In a simulated population of Marketplace insurance enrollees, for any given level of power, we find that a simple “full coverage after a deductible” reinsurance policy provides reductions in insurer risk comparable to the reductions offered by one-sided or two-sided risk corridors. Additionally, we find that one-sided risk corridors weakly dominate two-sided risk corridors by providing at least as much risk reduction for the same level of power. Our simulations also show that a simple reinsurance policy performs comparably to the set of risk-reducing policies being implemented in the Marketplaces consisting of partial coverage reinsurance and two-sided risk corridors.

It was not clear ex ante whether reinsurance or risk corridors would perform better. While an individual-level policy like reinsurance diminishes insurer incentives to constrain spending among individuals whose costs are likely to exceed the reinsurance threshold, such a policy has no effect on insurer incentives to limit spending among all other enrollees. Considering that the costs of only 0.45 percent of our sample exceed \$60,000 (the Marketplace reinsurance policy threshold), the average effect of reinsurance on power may be minimal in any given state of the world. Additionally, the 0.45 percent of the sample accounts for much of the variation in risk, suggesting that a policy such as reinsurance that reimburses insurers for those key individuals will substantially affect an insurer’s overall level of risk.

However, reinsurance, unlike risk corridors, affects insurer incentives for that small group of individuals in every state of the world. Risk corridors, on the other hand, affect insurer

²³The Marketplace policies have much larger effects on insurer risk for the 5,000 enrollee plan. These results are presented in the Online Appendix. Note, however, that the more relevant results regarding the risk-power trade-off are robust to plan size.

incentives for every enrollee but only in the rare case that the plan gets a bad draw from the cost distribution and incurs a large loss. Our simulations show that when measuring insurer incentives using the concept of power, plan-level one-sided risk corridor policies and individual-level reinsurance policies differ only slightly with respect to the risk-power trade-off we highlight in Section III.

The simulations also expose the shortcomings of two-sided risk corridors that finance the payments to plans with high costs through transfers from plans with low costs. While these two-sided policies seem “fair” they come with a cost: two-sided risk corridors limit insurer incentives to restrain utilization among both “winners” (plans with low costs) and “losers” (plans with high costs). In both cases, the price to the insurer of the marginal dollar spent on an enrollee’s health care is reduced by the two-sided risk corridor. With respect to the standard deviation measure of risk, this presents a trade-off between additional risk reduction and additional reductions in power. Our results show that the extra reduction in power exceeds the gains in risk reduction, generating our result that one-sided risk corridors outperform two-sided risk corridors.

These simulation results are subject to a few caveats. First, while our simulations accurately reflect the trade-off between insurer risk reduction and insurer incentives to constrain spending, they do not fully reflect the trade-off between insurer risk reduction and social welfare. Because the concept of power only describes the *incentives* of the insurer rather than actual behavior, it represents only an approximation of the efficiency consequences of these policies. A more informative analysis of risk-reducing policies would compare the effects of the policies on risk, holding constant the efficiency loss caused by the insurer’s behavioral response to the change in the insurer’s incentive to limit enrollee spending that we measure using power (i.e., insurer moral hazard). Insurer behavior in response to this incentive may be extreme or quite limited, depending on the cost to the insurer of efforts to limit enrollee spending. For example, one could imagine a situation where the cost to the insurer of modifying a high-cost enrollee’s total medical spending is extremely high. In such a situation, an insurer may never respond to a change in its incentive to constrain that enrollee’s spending, and the social cost of reimbursing an insurer for that enrollee’s spending would be zero. In some sense, then, power represents only the height of the Harberger triangle describing the welfare loss from reinsurance or risk corridors. While the other side of the triangle is necessary for a complete welfare analysis, any welfare loss associated with a given policy will be proportional to power, allowing us to use power to approximate the welfare consequences of these policies. Recovering the behavioral response of insurers to the shifts in incentives caused by these types of policies represents an important area for future research.

Second, while the simulations imply that when the regulator’s goal is to limit insurer risk for a given level of power reinsurance and one-sided risk corridors are comparable policies, the regulator may also have other goals. For example, the regulator may wish to limit the incentive for insurers to “cream-skim” low-cost enrollees. Reinsurance and risk corridors have dramatically different implications for insurers’ incentives to engage in this type of behavior. Reinsurance effectively redistributes costs from high-cost individuals to low-cost individuals, which dramatically weakens insurer incentives to select low-cost individuals

(Zhu et al. 2013; Geruso and McGuire 2014). Risk corridors, on the other hand, do not affect the heterogeneity of costs across the population of potential enrollees and thus do not affect the plan's incentives to select against sicker enrollees.²⁴ This suggests that when taking into account other objectives of the regulator, reinsurance may dominate risk corridors.

Risk corridors and reinsurance may also differ in how they affect an insurer's premium-setting incentives. If consumers are inertial in their choice of health plans, insurers may optimally choose to engage in "invest then harvest" strategies where they compete to be the "loss leader" in the first year of the operation of the market in order to gain market share (Ericson 2014). In subsequent years, insurers can then ratchet up the prices of their policies because consumer switching costs cause price sensitivity among incumbent enrollees to be quite low. With risk corridors, these strategies become more attractive to insurers because they limit the losses incurred from underpricing of policies during an individual's first year of enrollment. Two-sided risk corridors, on the other hand, may limit the attractiveness of these strategies by transferring money away from insurers earning large profits, reducing the surplus insurers can extract by ratcheting up prices in later years.

Finally, reinsurance and risk corridors may offer provider-sponsored insurance plans incentives to manipulate prices in order to extract reinsurance and risk corridor payments. Because for these provider-owned plans the prices their providers charge have no effect on the organization's bottom line, reimbursement that is tied to total reported cost (as provided by both reinsurance and risk corridors) provides an incentive for the organization to artificially inflate prices, resulting in potentially inefficient implicit transfers from non-provider-owned plans to provider-owned plans. It is possible that these incentives will be worse for risk corridors than for reinsurance because inflating prices will bring in much more revenue under risk corridors.

The implications of these unintended consequences of risk corridor policies represent important topics for future research. Additionally, while this paper demonstrates the potential value of risk corridors through a conceptual framework, empirical evidence on their effects in practice is needed. The effects of the risk corridor policies that have been implemented in Medicare Part D and in the early years of the ACA Marketplaces represent excellent opportunities for studying the effects of this policy in practice in the near future.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Research for this paper was supported by the National Institute of Mental Health (R01 MH094290, T32-019733).

²⁴Risk corridors can affect selection incentives if the contribution of low-cost individuals to the variance of the distribution is different from the contribution of high-cost individuals. If this is the case, then risk corridors do affect the incentives of risk-averse insurers to inefficiently select low-cost individuals because they will seek to do so for two reasons: (1) to shift the mean of the cost distribution to be as small as possible and (2) to decrease the variance. Because risk corridors affect the variance of the distribution and the effects of selection efforts on the variance of the distribution, they can affect selection incentives. However, selection incentives via the variance of the distribution are likely to be second-order to the incentives via the mean of the distribution.

References

- AcademyHealth. Profiles in Coverage: Arizona Healthcare Group. Robert Wood Johnson Foundation; 2007. http://www.academyhealth.org/files/publications/SCI_AZPiC.pdf [Accessed June 2, 2014]
- Aizcorbe, Ana; Liebman, Eli; Pack, Sarah; Cutler, David M.; Chernew, Michael E.; Rosen, Allison B. Measuring Health Care Costs of Individuals with Employer-Sponsored Health Insurance in the US: A Comparison of Survey and Claims Data. Bureau of Economic Analysis. 2012; 28:43–51. <http://www.bea.gov/papers/pdf/Measuring%20health%20care%20costs%20of%20individuals%20with%20employer-sponsored%20health%20insurance.pdf>.
- Arrow, Kenneth. Uncertainty and the Welfare Economics of Medical Care. American Economic Review. 1963; 53(5):941–73.
- Bovbjerg, Randall R.; Bowen Garrett, A.; Clemans-Cope, Lisa; Masi, Paul. State Coverage Initiatives: Reinsurance in State Health Reform. Reinsurance Institute Project of the Urban Institute, Robert Wood Johnson Foundation; 2008. http://www.academyhealth.org/files/publications/SCI_Reinsurance08.pdf [Accessed June 2, 2014]
- Boyersky, Victoria; Parke, Rob. The Medicare Shared Savings Program and the Pioneer Accountable Care Organizations. Milliman Healthcare Reform Briefing Paper. 2012. <http://us.milliman.com/uploadedFiles/insight/healthreform/medicare-shared-savings-program.pdf>
- Bulow, Jeremy. The Gaming of Pharmaceutical Patents. In: Jaffe, Adam B.; Lerner, Josh; Stern, Scott, editors. Innovation Policy and the Economy. Vol. 4. Cambridge, MA: MIT Press; 2004. p. 145–87.
- Centers for Medicare and Medicaid Services. Fast Facts – All Medicare Shared Savings Program ACOs. Centers for Medicare and Medicaid Services; 2013. <http://www.cms.gov/Medicare/Medicare-Fee-for-Service-Payment/sharedsavingsprogram/Downloads/All-Starts-MSSP-ACO.pdf> [Accessed June 2, 2014]
- Christofides, Stavros; Smith, Andrew D. DFA – The Value of Risk. Presented to the CAS 2001 DFA Forum; 2001. <http://www.casact.org/pubs/forum/01spforum/01spfl153.pdf>
- Dow, William H.; Fulton, Brent D.; Baicker, Katherine. Reinsurance for High Health Costs: Benefits, Limitations, and Alternatives. Forum for Health Econ Policy. 2010; 13(2):1–23.
- Einav, Liran; Finkelstein, Amy. Selection in Insurance Markets: Theory and Empirics in Pictures. Journal of Economic Perspectives. 2011; 25(1):115–38. [PubMed: 21595322]
- Ellis, Randall P.; McGuire, Thomas G. Provider Behavior under Prospective Reimbursement. Cost Sharing and Supply. Journal of Health Economics. 1986; 5(2):129–51. [PubMed: 10287223]
- Ericson, Keith MM. Consumer Inertia and Firm Pricing in the Medicare Part D Prescription Drug Insurance Marketplace. American Economic Journal: Economic Policy. 2014; 6(1):38–64.
- Fama, Eugene; Jensen, Michael. Agency Problems and Residual Claims. Journal of Law and Economics. 1983; 26(2):327–49.
- Geruso, Michael; McGuire, Thomas G. NBER Working Paper No. 20359. 2014. Tradeoffs in the Design of Health Plan Payment Systems: Fit, Power, and Balance.
- Glazer, Jacob; McGuire, Thomas G.; Shi, Julie. Risk Adjustment of Health Plan Payments to Correct Inefficient Plan Choice from Adverse Selection. NBER Working Paper No. 19998. 2013. <http://www.nber.org/papers/w19998>
- Gruber, Jonathan. Health Insurance and the Labor Market. NBER Working Paper No. 6762. 1998. <http://www.nber.org/papers/w6762>
- Hall, Brian J.; Liebman, Jeffrey B. Are CEOs Really Paid Like Bureaucrats? Quarterly Journal of Economics. 1998; 113(3):653–91.
- Handel, Benjamin R.; Hendel, Igal; Whinston, Michael D. Equilibria in Health Insurance Exchanges: Adverse Selection vs. Reclassification Risk. NBER Working Paper No. 19399. 2013. <http://www.nber.org/papers/w19399>
- Health and Human Services (HHS). [Accessed November 30, 2012] Patient Protection and Affordable Care Act; HHS Notice of Benefit and Payment Parameters for 2014. 2012. http://www.healthreformgps.org/wp-content/uploads/2012-29184_PI1.pdf
- Kahane, Yehuda. The Theory of Insurance Risk Premiums – A Re-examination in the Light of Recent Developments in Capital Market Theory. ASTIN Bulletin. 1979; 10:223–39.

- Laffont, Jean-Jacques; Tirole, Jean. A Theory of Incentives in Procurement and Regulation. Cambridge, MA: MIT Press; 1993.
- Layton, Timothy J. Working paper. 2014. Imperfect Risk Adjustment, Risk Preferences, and Sorting in Competitive Health Insurance Markets.
- Mahoney, Neale; Glen Weyl, E. Imperfect Competition in Selection Markets. NBER Working Paper No. 20411. 2014. <http://www.nber.org/papers/w20411>
- McGuire, Thomas G.; Glazer, Jacob; Newhouse, Joseph P.; Normand, Sharon-Lise; Shi, Julie; Sinaiko, Anna D.; Zuvekas, Samuel. Integrating Risk Adjustment and Enrollee Premiums in Health Plan Payment. *Journal of Health Economics*. 2013; 32(6):1263–77. [PubMed: 24308878]
- MedPAC. [Accessed June 2, 2014] Part D Payment System. 2012. http://www.medpac.gov/documents/MedPAC_Payment_Basics_12_PartD.pdf
- Newhouse, Joseph P. Reimbursing Health Plans and Health Providers: Efficiency in Production versus Selection. *Journal of Economic Literature*. 1996; 34(3):1236–63.
- Rothschild, Michael; Stiglitz, Joseph. Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *Quarterly Journal of Economics*. 1976; 90(4):629–49.
- Sappington, David EM.; Aydede, Sema K.; Dick, Andrew; Vogel, Bruce; Shenkman, Elizabeth. The Effects of Reinsurance in Financing Children's Health Care. *Inquiry*. 2006; 43(1):22–33.
- Shi, Julie. Working paper. 2013. Efficiency in Plan Choice with Risk Adjustment and Premium Discrimination in Health Insurance Marketplaces.
- Sing, Merrile; Banthin, Jessica S.; Selden, Thomas M.; Cowan, Cathy A.; Keehan, Sean P. Reconciling Medical Expenditure Estimates from the MEPS and NHEA. *Health Care Financing Review*. 2002; 28:25–40. [PubMed: 17290666]
- Swartz, Katherine. Reinsuring Health: Why More Middle-Class People Are Uninsured and What Government Can Do. New York: Russell Sage Foundation; 2006.
- Ross, Winkleman; Peper, Julie; Holland, Patrick; Mehmud, Syed; Woolman, James. Analysis of HHS Proposed Rules on Reinsurance, Risk Corridors, and Risk Adjustment. State Health Reform Assistance Network, Robert Wood Johnson Foundation, Wakely Consulting Group; Boston, MA: 2012. <http://www.statenetwork.org/wp-content/uploads/2012/04/State-Network-Wakely-Analysis-of-HHS-Final-Rules-On-Reinsurance-Risk-Corridors-And-Risk-Adjustment.pdf>
- Zeckhauser, Richard J. Medical Insurance: A Case Study of the Trade-off between Risk Spreading and Appropriate Incentives. *Journal of Economics Theory*. 1970; 2(1):10–26.
- Zhu, Julie; Layton, Timothy J.; Sinaiko, Anna D.; McGuire, Thomas G. The Power of Reinsurance in Health Insurance Marketplaces to Improve the Fit of the Payment System and Reduce Incentives for Adverse Selection. *Inquiry*. 2013; 50(4):255–75. [PubMed: 24996751]
- Zuvekas, Samuel H.; Olin, Gary L. Accuracy of Medicare Expenditures in the Medical Expenditure Panel Survey. *Inquiry*. 2009; 46(1):92–108. [PubMed: 19489486]

FIGURE 1. Insurer's distribution of expected costs

Notes: This figure illustrates the expected cost distribution faced by insurers. Possible realizations of per capita costs are on the x-axis, and the density is on the y-axis. The expected cost is the mean of this distribution. We assume that prices are set such that any outcome to the left of the mean (lower than expected costs) results in positive profits and any outcome to the right of the mean (higher than expected costs) results in losses.

FIGURE 2. Results of policy simulations (insurer risk)

Notes: Each figure shows the results of simulations of insurer risk under a variety of reinsurance and risk corridor policies. All reinsurance policies reimburse the insurer for 100% of an individual's costs beyond a given cutoff. This cutoff is shown on the top x-axis, covering the relevant range of \$0–\$350,000. All one-sided (two-sided) risk corridor policies reimburse the insurer for 100% of plan costs exceeding a fixed percentage of the mean of the cost distribution, and are funded through a uniform per capita actuarially fair “risk corridor premium” (through symmetric transfers away from plans whose costs are below a fixed percentage of the mean of the cost distribution). The fixed percentage is shown on the bottom x-axis, covering the relevant range of 1–1.05.

FIGURE 3. Results of policy simulations (power)

Notes: The figure shows the results of simulations of power under a variety of reinsurance and risk corridor policies. All reinsurance policies reimburse the insurer for 100% of an individual's costs beyond a given cutoff. This cutoff is shown on the top x-axis, covering the relevant range of \$0–\$350,000. All one-sided (two-sided) risk corridor policies reimburse the insurer for 100% of plan costs exceeding a fixed percentage of the mean of the cost distribution, and are funded through a uniform per capita actuarially fair “risk corridor premium” (through symmetric transfers away from plans whose costs are below a fixed percentage of the mean of the cost distribution). The fixed percentage is shown on the bottom x-axis, covering the relevant range of 1–1.05.

FIGURE 4. Joint distribution of risk and power from policy simulations

Notes: Each figure shows the joint distribution of power and risk generated from simulations of a range “full coverage after a deductible” reinsurance and risk corridor policies with varying cutoffs for reimbursement from the regulator. Risk is on the y-axis and power is on the x-axis. For any given level of power, the lines show the level of risk achieved by a reinsurance or risk corridor policy generating that level of power. For reinsurance, we simulate policies reimbursing plans for 100% of an individual’s costs above some threshold ranging from \$0 to \$375,000. For risk corridors, we simulate policies reimbursing plans for all plan-level costs exceeding a fixed percentage of the mean of the cost distribution. We allow the percentage to range from 100.1% to 130%. Under two-sided risk corridors, plans with low costs transfer money to plans with high costs. One-sided risk corridor transfers are funded through a uniform per capita actuarially fair risk corridor premium.

TABLE 1

Summary statistics of the full population

<i>N</i> = 76647			
Demographics		Income	
Age 0–18	0.16	Family income	\$69,103.88
Age 19–34	0.33	Medical	
Age 35–44	0.19	Total annual expenditures	\$3,103.48
Age 45–64	0.32	Insurance status	
Male	0.51	Uninsured	0.46
Married	0.49	Non-group	0.04
Race		Expensive ESI	0.06
White, non-Hispanic	0.50	Self employed	0.01
Black, non-Hispanic	0.14	Medicaid	0.10
Hispanic	0.28	Small group ESI	0.27
Other	0.08	Health status	
Education		Excellent	0.32
Less than high school	0.19	Very good	0.32
High school	0.29	Good	0.27
Some college	0.15	Fair	0.08
College degree	0.25	Poor	0.02
Employment status			
Continuously employed	0.64		
Continuously unemployed	0.11		

Notes: Values in table are means unless otherwise noted. Statistics calculated for Marketplace-eligible population constructed from panels 9–14 of the MEPS. We select adult, non-elderly individuals in households earning at least 138% of FPL and children in households with income of at least 205% of FPL. We include households where an adult is ever uninsured, a holder of a non-group insurance policy, self-employed, employed by a small employer, or paying an out-of-pocket premium for their employer-sponsored health insurance (ESI) plan that is deemed to be unaffordable. If an individual meets the selection criteria in at least one of the two survey years, she is part of the sample. The dataset comprises 44,210 “Marketplace-eligible” individuals, 11,773 of whom have only one year of data and 32,437 of whom have two years of data, generating a total sample size of 76,647 person-years.

TABLE 2

Insurer risk under proposed policies

Risk measure	Base case	Reinsurance only	Two-sided risk corridor only	Reinsurance and two-sided risk corridor
Standard deviation	\$61.02	\$44.48	\$57.98	\$44.03
95th percentile of cost distn	\$3,424.51	\$3,397.10	\$3,424.51	\$3,397.10
99th percentile of cost distn	\$3,471.75	\$3,427.85	\$3,448.27	\$3,426.31
Power	1	0.89	0.95	0.88

Notes: Each column represents a different set of policies and each row represents a moment of the insurer's cost distribution under those policies. The cost distribution in the base case is approximated by taking 5,000 draws of 20,000 individuals from the population of Marketplace-eligible individuals we create from MEPS data. The cost distributions under reinsurance, risk corridors, and both reinsurance and risk corridors are approximated by starting with the base case distribution and then calculating the reinsurance and risk corridor transfers and fees according to the parameters proposed for implementation in the Marketplaces.