NBER WORKING PAPER SERIES

CATASTROPHE BONDS, REINSURANCE, AND THE OPTIMAL COLLATERALIZATION OF RISK-TRANSFER

Darius Lakdawalla
George Zanjani

Working Paper 12742
http://www.nber.org/papers/w12742

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2006

The views expressed in this paper are those of the authors, and do not necessarily reflect the views of the Federal Reserve Bank of New York, the Federal Reserve System, the RAND Corporation, or the RAND Center for Terrorism Risk Management and Policy. We thank Robert Craig, Ken Garbade, Robert M. Hall, Thomas Holzheu, Stewart C. Myers, and seminar participants at the 2006 NBER Insurance Project Workshop and the 2006 Risk Theory Seminar for helpful comments and discussions. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Catastrophe Bonds, Reinsurance, and the Optimal Collateralization of Risk-Transfer
Darius Lakdawalla and George Zanjani
NBER Working Paper No. 12742
December 2006
JEL No. G11,G22

ABSTRACT

Catastrophe bonds feature full collateralization of the underlying risk transfer, and thus abandon the insurance principle of economizing on collateral through diversification. We examine the theoretical foundations beneath this paradox, finding that fully collateralized instruments have important uses in a risk transfer market when insurers cannot contract completely over the division of assets in the event of insolvency, and, more generally, cannot write contracts with a full menu of state-contingent payments. In this environment, insureds have different levels of exposure to an insurer's default. When contracting constraints limit the insurer's ability to smooth out such differences, catastrophe bonds can be used to deliver coverage to those most exposed to default. We demonstrate how catastrophe bonds can improve welfare in this way by mitigating differences in default exposure, which arise with: (1) contractual incompleteness, and (2) heterogeneity among insureds, which undermines the efficiency of the mechanical pro rata division of assets that takes place in the event of insurer insolvency.

Darius Lakdawalla
The RAND Corporation
1776 Main Street
Santa Monica, CA 90407-2138
and NBER
darius@rand.org

George Zanjani
Capital Markets Function
Federal Reserve Bank of New York
33 Liberty St.
New York, NY 10045
george.zanjani@ny.frb.org

1 Introduction

Recent disaster experience has produced a flurry of economic inquiry into catastrophe insurance markets. Especially puzzling is the apparent incompleteness of catastrophe risk transfer: The price of risk transfer seems high, risk is not spread evenly among insurers in the manner suggested by Borch's [1] groundbreaking theoretical result, and, in stark contrast to Arrow's well-known characterization of optimal insurance contracts, reinsurance consumers do not purchase coverage for high layers of risk. Froot [11] documents these puzzles and fingers various market imperfections as possible explanations.

Some view the catastrophe bond as a promising vehicle for overcoming imperfections in the reinsurance market. In principle, the catastrophe bond opens a direct channel for catastrophe risk to flow to the capital markets, sidestepping frictions present in the reinsurance market and connecting those who need protection with well-funded investors eager to provide it. On the other hand, others are skeptical that catastrophe securitization will be a panacea. Bouriaux and Scott [2] argue that securitization terms are unlikely to be attractive to buyers of terrorism coverage and note that the record of risk-linked capital market instruments has not been encouraging. Indeed, catastrophe bond issuance to date has been relatively limited, even in the aftermath of events that were expected to stimulate issuance. It is too early to write an epitaph for the catastrophe bond, but the experience to date raises questions about its theoretical foundations and its likely future role.

On closer inspection, the catastrophe bond seems both paradoxical and primitive. Its current form features full collateralization and links principal forfeiture only to specific risks, thereby retreating from the fundamental, time-tested concept of diversification that allows insurers to protect insured value far in excess of the actual assets held as collateral. In a world where frictional costs (e.g., due to taxes, regulations, or moral hazard) make capital costly to hold, diversification lowers the cost of insurance. Viewed in this light, a fully collateralized instrument seems an unlikely competitor to traditional reinsurance products.[1]

This paper examines this issue by developing a theory of risk collateralization. Specifically, we study the efficient division of risk-bearing assets between reinsurance company assets and catastrophe bond principal (both of which can be used to "collateralize" promises to indemnify consumers). In a narrow sense, the analysis supports the skeptical intuition outlined above. When reinsurance companies can write any type of contract with their insureds and frictional costs are identical for catastrophe bonds and reinsurer assets, catastrophe bonds are at best redundant, and at worst welfare-reducing. If the insurer has complete freedom to vary indemnity payments to consumers in every state of the world, it can engineer any possible menu of pay-outs through its own contracts: Catastrophe bonds add nothing in the absence of contracting constraints.

However, reinsurers and insurers *do* face contracting constraints in practice. Typical contracts promise an indemnity payment to a policyholder who has suffered a covered loss but do not specify rules ex ante for who gets what in the event of insurer failure. Instead, the division of assets under bankruptcy is determined by insurance receivership laws, and insurers either cannot or do not attempt to contractually specify how assets will be divided in the event of bankruptcy. As a result, assets of failed companies are distributed according

[1]Niehaus [21] observes this paradox on Page 593.

to inflexible and potentially inefficient mechanical rules.

Constraints on company asset distribution under default open up a role for catastrophe bonds (and other highly collateralized vehicles, such as reinsurance "side-cars"), even if there are no differences in the frictional costs. When insureds are homogeneous and risk exposures are binary (loss or no loss), optimal insurance contracts are similar, and simple rules—e.g., a *pro rata* rule that pays all claimants at the same rate on the dollar in the event of insurer bankruptcy—can perform well. Heterogeneity, however, exposes the shortcomings of insurance contracts in the presence of pro rata rules, which may misallocate assets in the bankruptcy state.[2] Pro rata allocations can be suboptimal when some insureds are more concerned about the bankruptcy state than others, and this will generally be the case. Reinsurance buyers hold policies of differing quality even when purchasing these policies from the same reinsurer. Some are more exposed to default than others, and those that have greater exposure to bankruptcy risk may desire greater collateralization of their potential claims than can be provided under pro rata rules. This need opens up a role for catastrophe bonds and other related instruments in the risk transfer market, which can vary the degree of collateralization of coverage for specific insureds.

Thus, catastrophe bonds can improve welfare when reinsurers face constraints on the distribution of assets in bankruptcy, *and* when they must insure a heterogeneous group of risks. Catastrophe bonds can smooth out allocations made ragged by the risk of bankruptcy. Put differently, reinsurance capital may weakly dominate the catastrophe bond in terms of raising *average* policy quality, but such capital can be rendered a blunt instrument by bankruptcy laws: Catastrophe bonds can improve welfare for those insureds most exposed to bankruptcy risk.

The paper is laid out as follows. Section 2 provides some background and context on catastrophe bonds. Section 3 then develops a concrete two-consumer example to illustrate the intuition behind our results. Section 4 develops our results formally in the context of a social planning problem with N consumers. Section 5 discusses other strategies for protecting consumers against default and interprets them in the context of the model. In particular, it considers how collateralization clauses in reinsurance policies influence the priority of claimants under bankruptcy, and the extent to which such clauses substitute for fully collateralized instruments such as catastrophe bonds. Section 6 concludes.

2 Background and Motivation

While structures vary, a typical catastrophe bond transaction involves a special purpose vehicle (SPV). The SPV sells securities (catastrophe bonds) to investors, and the proceeds from the sale are deposited in trust and invested. The SPV then provides reinsurance to a ceding insurer or reinsurer, who pays a premium in exchange. The premium, as well as income earned on the trust investments, funds interest payments to investors. If a contractually-defined trigger event occurs, part or all of the bond principal is forfeited to the

[2]Mahul and Wright [17] also note the inefficiency of *pro rata* rules in the context of a model with identical consumers but generalized loss distributions.

ceding company; if no event occurs, the principal is returned to investors.[3]

The catastrophe bond market is still evolving and, in particular, moving toward higher layers of risk. While the first catastrophe bonds linked forfeiture of principal to the issuer's actual losses (an *indemnity trigger*), triggers linking forfeiture of principal to industry losses, model output, or to specific parameters of the disaster (e.g., the strength of an earthquake centered in a certain geographic region) have grown in popularity. Some deals feature multiple event triggers—requiring two or more major disasters within a short time period to trigger principal forfeiture (see Woo [22]).

How important is the catastrophe bond market? Catastrophe bond issuance in 2005 amounted to a record $2 billion, with outstanding principal in the neighborhood of $5 billion.[4] With ongoing growth in catastrophe bond issuance, as well as the emergence of reinsurance side-cars (which will be discussed later), 2006 is set to be another record year for alternative risk transfer in the catastrophe market. On the other hand, catastrophe bonds currently provide only a tiny fraction of reinsurance capacity. Despite being around for more than a decade, outstanding catastrophe bond principal amounted to less than 2% of global reinsurance capital at year-end 2005.[5] Relative to the giddy expectations of the 1990's, the volume of catastrophe securitization has been disappointing.

Nevertheless, while issuance so far has fallen short of optimistic projections, we argue that catastrophe bonds do serve a well-defined economic role in the risk transfer market—a role deriving from the full collateralization underlying the instrument. Full collateralization allows catastrophe bonds to be useful in cases where traditional risk transfer (e.g., through reinsurance policies) is subject to significant risk of counterparty default. This role could expand if frictional costs associated with catastrophe bonds (e.g., issuance costs, secondary market illiquidity) fall significantly in relation to the frictional costs associated with reinsurance equity. However, unless frictional costs become negligible, full collateralization is likely to limit the role of catastrophe bonds in circumstances where diversification opportunities make partially collateralized instruments attractive.

To understand the role of customer-specific collateral, we must move beyond existing theory on insurance pricing. The theory of the insurance firm has made a great deal of progress in understanding the joint determination of multiple line pricing, capital allocation, and the firm's overall default risk (see, e.g., Cummins et al. [6], Myers and Read [20], and Zanjani [23]). However, most models consider the default risk of the firm *as a whole.* Less progress has been made in studying differences across a firm's policyholders in their exposure to default.

Differences across policyholders, though, are central to the value of catastrophe bonds. If the object of interest is a single default-related financial target for the company as a whole—such as the expected policyholder deficit per dollar of liabilities—a dollar held in the form of a catastrophe bond cannot possibly be preferable to one held as company equity.

[3]For more information on the structure of insurance-linked securities, see, for example, Canabarro, E., Finkemeier, M., Anderson, R.R., and Bendimerad, F. (1998), "Analyzing Insurance-Linked Securities," Goldman Sachs & Co. Quantitative Research, October, 1998.

[4]Source: *The Catastrophe Bond Market at Year-End 2005: Ripple Effects from Record Storms,* MMC Securities. Figures include only publicly disclosed transactions.

[5]Standard and Poor's *Global Reinsurance Highlights, 2006 edition* tallied over $300 billion in total adjusted shareholder funds for the industry at year-end 2005.

Since the dollar held as equity will be available in all states of the world, it will be available to pay for all of the losses that will be covered by a catastrophe bond and some losses that are not covered by the catastrophe bond.

To understand how catastrophe bonds can be used in the absence of compelling frictional cost differences, we must move beyond thinking of a single default-related financial target for the insurance company. Instead, we must think about the company's policies as having varying levels of quality, corresponding to varying levels of exposure to default, and how catastrophe bonds and equity have distributional consequences for recoveries by different policyholder groups in states of default.

3 A Simple Example

In the context of a simple two-consumer example, we illustrate how the role for catastrophe bonds depends on the presence of: (1) nonzero bankruptcy risk for the insurer; (2) contracting constraints that prevent the insurer from optimally allocating claims payments in the bankruptcy state; and (3) heterogeneity across consumers, such that one consumer faces greater exposure to insurer bankruptcy risk.

Consider the case of two consumers, named A and B. Consumer A faces a 10% chance of losing \$100, while Consumer B faces a 1% chance of losing \$100. An insurer issues simple contracts to indemnify the consumer, fully or partially, in the event of a loss. In the bankruptcy state (where claims exceed insurer assets), claims payments are allocated according to a mechanical rule by dividing assets on a pro-rata basis, according to the claims made by the insureds.[6]

Suppose we have \$150 in assets. How should we allocate them? Consider first the case where we use all \$150 to fund an insurance company, which issues a \$100 limit insurance policy to A and a \$100 limit policy to B. Expected claims in this example equal:

$$10\% * \$100 + 1\% * \$100 = \$11 \tag{1}$$

The insurer is able to pay all claims in full except when both consumers suffer a loss; in that event, the insurer pays out all \$150 of its assets but declares bankruptcy. Therefore, expected claims payments equal:

$$10\% * 99\%(\$100) + 1\% * 90\%(\$100) + 10\% * 1\%(\$150) = \$10.95 \tag{2}$$

Overall, the insurer pays $\$\frac{10.95}{11}$, or better than 99 cents, on the dollar. However, the two consumers are unevenly exposed to default. Consumer B ends up being much more exposed to bankruptcy risk on a per dollar basis, because she faces a higher relative likelihood of suffering a loss in the state of the world where the other consumer *also* suffers a loss.

Specifically, Consumer A expects to lodge \$10 worth of claims and to receive payments of:

$$10\% * 99\% * \$100 + 10\% * 1\% * \$75 = \$9.975 \tag{3}$$

[6] The exact form of the mechanical rule is less relevant than the presence of contracting constraints in the bankruptcy state.

On the other hand, Consumer B expects to lodge \$1.00 worth of claims, but receive

$$1\% * 90\% * \$100 + 1\% * 10\% * \$75 = \$0.975 \tag{4}$$

Thus, Consumer A receives 99.975 cents on the dollar, while Consumer B receives only 97.5.

Consumer A is better insured than Consumer B, and the social planner, depending on the objective function, might want to consider redistributing coverage from Consumer A to Consumer B. One way of accomplishing this is to redeploy some of our assets in the form of a catastrophe bond tied to Consumer B. Suppose we now spend \$100 funding the insurance company, which sells a \$100 limit insurance policy to Consumer A and a \$50 limit policy to Consumer B. We then use the remaining \$50 on a catastrophe bond payable to Consumer B in the event of a loss.

Consumer A still expects to lodge \$10 worth of claims, but now receives payments of:

$$10\% * 99\% * \$100 + 10\% * 1\% * (\frac{100}{150} * 100) = \$9.967 \tag{5}$$

On the other hand, Consumer B now expects to lodge \$0.50 worth of claims with the insurance company, but now also is entitled to receive \$50 of catastrophe bond principal in the event of a loss:

$$1\% * 90\% * (\$50 + \$50) + 1\% * 10\% * (\frac{50}{150} * \$50 + \$50) = \$0.983 \tag{6}$$

In other words, the recovery differential narrows. Consumer A now receives 99.67 cents of relief per dollar of loss, a slightly worse rate than before. Consumer B now receives a bit more—98.3 cents—once the catastrophe bond principal is considered.

In this example, using the catastrophe bond instead of a full reinsurance solution effectively transfers coverage from one consumer to the other. The transfer occurs only when the reinsurer defaults: We have sufficient assets to fully indemnify both consumers *except* when both experience a loss, and the catastrophe bond allows us to affect the distribution of indemnification in that unfortunate state of the world. Of course, the question of whether or not this redistribution is desirable depends on particulars such as preferences—but the general point is that the allocation of assets to consumers in the bankruptcy state may be suboptimal in a pure reinsurance solution, and the catastrophe bond is one way of securing the interests of one consumer over the other.

The presence of contracting constraints, the risk of bankruptcy, and the presence of consumer heterogeneity all play key roles in driving this result. If an insurer is able to write complex contracts that vary indemnification across all states of the world, it can replicate the payout structure of a catastrophe bond without using the bond itself. For instance, in the example above, we could replicate the payoffs involved under the second approach (using the catastrophe bond) simply by capitalizing the insurer with \$150 and issuing polices offering full \$100 indemnification except in the case where both consumers had losses, in which case Consumer A would receive \$66.67 and Consumer B would receive \$83.33.[7]

[7] Note that if we allowed policy limits to exceed insurer assets, this would allow insurers to influence the division of resources in the bankruptcy state. However, this is a blunt instrument for resource allocation that cannot generally replicate the payouts of catastrophe bonds. For example, with more than two consumers, insurer bankruptcy is not perfectly correlated with the loss experience of any one consumer, because there are many possible loss configurations that trigger bankruptcy. Nevertheless, in the general theory developed in Section 4, we place no constraints on the choice of policy limits.

Contracting constraints that prevent the insurer from specifying complicated priority rules under bankruptcy are necessary to preclude this possibility. Heterogeneity also plays an important role by rendering mechanical bankruptcy rules inefficient. If Consumers A and B were identical, an equal pro rata division of resources in the bankruptcy state would be optimal, and neither consumer would be any more exposed to default risk.

This example shows how catastrophe bonds can be used to improve social welfare by redistributing coverage among consumers in "unfortunate" states of the world, but it falls short of illustrating other aspects of the general trade-off between catastrophe bonds and reinsurance. Earlier, we emphasized the costliness of fully collateralized catastrophe bonds, relative to less than fully collateralized insurance. Yet in this example, there is no disadvantage to "sequestering" capital in the form of a catastrophe bond since we make full use of the collateral assets. In more general versions of the problem, one of the important drawbacks associated with catastrophe bonds is that the assets dedicated to catastrophe bond principal for one consumer will not be available to pay losses experienced by others. In the results that follow, we show that the three conditions identified above are necessary for catastrophe bond issuance to be useful, but not sufficient—any benefits associated with catastrophe bond issuance may fail to outweigh the inefficiencies associated with full collateralization.

4 Theory

The basic approach borrows from Borch's analysis of optimal risk sharing among many consumers: Instead of modelling individual behavior, we study the social planning problem and its solutions.

Consider a world with N consumers. Consumer i is endowed with initial wealth W_i and faces the risk of experiencing a loss of fixed size—denoted by L_i. To characterize the possible states of the world, we define a row vector $\mathbf{x}$ of length N, with the elements all taking a value of zero or one: $\mathbf{x}(i) = 1$ means that consumer i experienced a loss, while $\mathbf{x}(i) = 0$ means that she did not. Let Ω denote the set of all such vectors of length N with the elements taking values of one or zero. Each element of Ω corresponds to a complete description of one possible state of the world. The entire set Ω contains all possible such states. The following set definitions are useful:

$$\Omega^i = \{\mathbf{x} : \mathbf{x}(i) = 1\},$$

the set of all states in which agent i suffers a loss, and

$$\Gamma(\mathbf{x}) = \{i : x(i) = 1\},$$

the set of all agents that suffer a loss in state $\mathbf{x}$.
Thus, using this notation, we may describe the probability of loss faced by consumer i as:

$$p_i = \sum_{\mathbf{x} \in \Omega^i} \Pr(\mathbf{x})$$

There are two risk transfer technologies available to insure against losses. First, we can set up an insurance company and issue insurance policies to consumers, collateralized by the assets of the company. Second, we can issue a risk-linked security on behalf of a consumer (i.e., a catastrophe bond) that pays off in the event that the consumer experiences a loss.

The insurance company is formed with assets of A. Throughout our discussion, we think of "assets" as all the resources the insurer can use to pay claims. Therefore, it includes both capital paid in by investors and premiums paid in by consumers. For our purposes, the key issue is whether or not the dollar is available for claims payment—not how it would be treated by accounting conventions. In the event of a loss (or losses), the consumers can draw on the assets to pay claims. When assets exceed claims, what remains after claims are paid reverts to investors. On the other hand, if claims exceed assets, the company defaults, and claimants are assumed to be paid according to a pro rata rule—everyone receives the same rate of recovery per dollar of claim.

Each consumer can separately issue a catastrophe bond[8] to investors. The principal of the bond is forfeited to the consumer in the event of a loss, but not otherwise. Let B_i be the bond issuance of consumer i. We simplify matters by assuming *indemnity* triggers—where principal forfeiture is linked directly to issuer loss experience. Hence, we avoid the complexities of optimal trigger design (see Doherty and Mahul [7]) and the problem of basis risk. We do not directly model these and other costs associated with asymmetric information,[9] but such costs can be thought of as being embedded in the frictional costs associated with catastrophe bond principal described below.

What do insurance policies and catastrophe bonds cost? The cost of risk transfer can be decomposed into 1) fair ex ante compensation for claims expected to be paid under the risk transfer agreement and 2) frictional costs associated with establishing and maintaining the risk transfer scheme. We start simply by assuming that the cost of risk transfer amounts to the expected value of claims plus a frictional cost proportional to the amount of collateral used in the risk transfer scheme (i.e., the amount of assets used in the insurance company, or the amount of catastrophe bond principal used). We then generalize the results to the case where fair compensation for expected claims reflects the states of the world where the claims occur: That is, we price risk transfer according to how claims are expected to relate to the expected distribution of returns on other assets available in the broader capital markets.

We start "simply" because it turns out that our results are driven entirely by frictional costs. In the absence of frictional costs, all consumers will be fully insured, and it is irrelevant how risk transfer technologies are combined in providing this full insurance: Frictional costs provide a motivation to economize on capital in the process of collateralizing risk-transfer. Accordingly, we start with a model where insurance risks are "zero beta," but where frictional capital costs lead to limited risk transfer. We then show that these results hold even when insurance risks correlate in some way with capital market returns (and are priced accordingly), but that non-zero frictional costs are key to our findings.

[8] We refer to "catastrophe bonds" because of their familiarity, but the analysis that follows obviously applies to other highly collateralized instruments used in risk transfer—such as collateralized reinsurance policies and "side-cars." These and other risk transfer alternatives will be discussed in Section 5.

[9] See Brandts and Laux [3] for a theoretical justification for the catastrophe bond market based on asymmetric information between insurers and reinsurers.

4.1 Optimal Collateralization with Frictional Costs

Frictional costs are imagined here as deriving from agency costs, taxes, liquidity costs, or other frictions. Each dollar of assets held in the insurance company results in per unit frictional costs of δ_A. Each dollar of catastrophe bond principal raised has the frictional cost δ_B.

Consumers must pay for all frictional costs and fair compensation for recoveries expected from the risk transfer. Denote the portion of this total risk transfer cost allocated to consumer i as c_i.

Insurance policies are represented as simple promises of indemnification: The insurer promises to pay I_i in the event that consumer i experiences a loss. We place no constraints on the promised indemnity: It may be less than, equal to, or greater than the prospective loss. However, contracting constraints come into play in the sense that we do not allow the insurer to offer a schedule of promised indemnification, with the amount contingent on the loss experiences of other insureds. If the insurer is able to pay, it pays in full; if not, it defaults, and all claims are paid at the same rate on the dollar. The example of Section 3 suggests that relaxing this contracting constraint will obviate roles for catastrophe bonds or other fully collateralized instruments. In the Appendix (Section A), we verify that this is in fact the case: When frictional costs are identical $(\delta_A \equiv \delta_B)$, allowing the social planner to arbitrarily vary the indemnification promised in each insurance policy eliminates any potential role for catastrophe bonds.

We can now define utility for consumer i (according to the usual Von Neumann-Morgenstern assumptions) as:

$$EU_i = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) U_i \left(W_i - L_i + f_{\mathbf{x}} I_i + B_i - c_i\right) + \sum_{\mathbf{x}\notin\Omega^i} \Pr(\mathbf{x}) U_i \left(W_i - c_i\right), \tag{7}$$

where $f_{\mathbf{x}}$ represents the proportion of the indemnity payment actually paid in state $\mathbf{x}$.

The social planning problem can now be written as:

$$\max_{A,\{B_i\},\{c_i\},\{I_i\},\{f_{\mathbf{x}}\}} V = \sum_i EU_i \tag{8}$$

subject to:

$$[\mu] \quad : \quad \sum c_i \geq \delta_A A + \delta_B \sum_i B_i + \sum_{\mathbf{x}\in\Omega} \Pr(\mathbf{x}) \left(\sum_{i\in\Gamma(\mathbf{x})} f_{\mathbf{x}} I_i + \sum_{i\in\Gamma(\mathbf{x})} B_i \right) \tag{9}$$

$$[\lambda_{\mathbf{x}}] \quad : \quad f_{\mathbf{x}} \sum_{i\in\Gamma(\mathbf{x})} I_i \leq A, \forall \mathbf{x} \tag{10}$$

$$[\phi_{\mathbf{x}}] \quad : \quad f_{\mathbf{x}} \leq 1, \forall \mathbf{x} \tag{11}$$

and subject to non-negativity constraints on catastrophe bond principal and policy limits. Constraint 9 ensures that consumers' total payments for risk-transfer instruments ($\sum c_i$) cover the frictional costs of capital and expected losses. Constraint 10 ensures that the insurer always has enough assets on hand to cover *actual* (as opposed to promised) liabilities. Finally, constraint 11 precludes the insurer from ever paying out more than the policy limit.

Any difference in frictional costs (e.g., $\delta_A \neq \delta_B$) will obviously create a potential advantage for one of the technologies, but we will start by considering the case where

$$\delta_A \equiv \delta_B \equiv \delta.$$

Thus, we start by studying how the nature of preferences and risk affect the optimal mix of the two risk transfer technologies. The optimality conditions are derived in the Appendix (Section B.1), as is the following marginal condition for catastrophe bond issuance (where we use the notation $U_i^{\mathbf{x}}$ to denote the utility of consumer i in state $\mathbf{x}$):

$$\boxed{R_i = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x})\,[1-f_{\mathbf{x}}] \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \sum_{j\in\Gamma(\mathbf{x})} w_j^{\mathbf{x}} \frac{\partial U_j^{\mathbf{x}}}{\partial W} \right) - \sum_{\mathbf{x}\notin\Omega^i} \lambda_{\mathbf{x}} \leq 0} \tag{12}$$

where

$$w_j^{\mathbf{x}} = \frac{I_j}{\sum_{j\in\Gamma(\mathbf{x})} I_j}.$$

R_i is the "marginal return on catastrophe bonds"—the increase in utility associated with the first-dollar of bond issuance. $R_i \leq 0$ if and only if catastrophe bonds cannot improve on a reinsurance-only equilibrium. Specifically, if R_i is negative, this means that catastrophe bond issuance was not useful (optimal) for consumer i—or, in other words, that $B_i^* = 0$. Study of (12) reveals several important characteristics of optimal catastrophe bond use in the social planning problem.

First, a catastrophe bond's potential to enhance the welfare of the issuing consumer is intimately linked to the presence of default risk. If consumer i does not face any risk of default (i.e., $f_{\mathbf{x}} = 1$ for all $\mathbf{x} \in \Omega^i$), catastrophe bond issuance will not be useful for that consumer.[10]

Second, assuming consumer i is confronted with default, catastrophe bond issuance on behalf of that consumer has the *potential* to be useful only if her marginal valuation of consumption in the states where the company defaults on her claim exceeds the average valuation of other consumers who lose in those same states:

$$\sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x})\,[1-f_{\mathbf{x}}] \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \sum_{j\in\Gamma(\mathbf{x})} w_j^{\mathbf{x}} \frac{\partial U_j^{\mathbf{x}}}{\partial W} \right) > 0.$$

Intuitively, it will not make sense to dedicate collateral to consumer i if that collateral is actually worth more to her companions in states of default. If her companions value that collateral more highly on average than she does, it will be more efficient to add that

[10]This is equivalent to saying that $R_i < 0$, except in solutions where the insurance company never defaults on any contract. If the insurer never defaults, $R_i = 0$, implying that cat bonds could figure in a solution. However, as shown in the Appendix (Section B.2), any such solution would not be unique: Without default, any solution with catastrophe bond principal can be matched by a solution without catastrophe bond principal.

collateral to the insurance company or to issue catastrophe bonds on behalf of the high valuation consumers.

Finally, the value of catastrophe bond issuance for consumer i also depends on the extent of diversification possibilities, as captured in:

$$\sum_{\mathbf{x} \notin \Omega^i} \lambda_{\mathbf{x}}.$$

If that term is positive, it means those diversification possibilities still exist. In other words, the company is defaulting in states of the world where consumer i does not experience a loss, and, thus, there are other consumers who would enjoy benefits from increasing the capitalization of the insurance company. While this does not preclude the solution from featuring catastrophe bond issuance on behalf of consumer i, it makes it more difficult for the catastrophe bond to be the preferable instrument for addressing the the risk transfer needs of the consumer in question: Any benefits obtained by sequestering collateral on behalf of consumer i (and thus shielding the assets from consumers who place lower valuations on additional coverage in those states where consumer i is exposed to default) must be weighed against the cost of preventing consumers exposed to default in other states of the world from accessing that collateral.

The importance of differences across consumers is highlighted by the case where consumers are homogenous, or ex ante identical. The Appendix explores this case in detail, showing that catastrophe bonds will not be useful in this case. The result (shown in the Appendix - Section B.3) can be understood by noting that, assuming that symmetric solutions apply under homogeneity, the marginal utilities of consumers who lose will be equivalent in each state. So Equation 12 reduces to:

$$R_i = -\sum_{x \notin \Omega^i} \lambda_{\mathbf{x}} \leq 0$$

In other words, cat bond issuance will be strictly suboptimal unless diversification possibilities have been completely exhausted. More precisely, cat bond issuance for consumer i will be suboptimal if there is a *single* state of the world where additional insurance capital will benefit someone *other* than consumer i. In the case of homogeneity, this can only happen if all consumers enjoy full indemnification except in the state where everyone experiences a loss. If the social planner finds it desirable to indemnify consumers to this extent, she will be indifferent between cat bonds and insurance policies as a means of providing additional coverage in the N-loss state.[11]

Thus, the N-consumer case under homogeneity exposes the extreme disadvantage of cat bonds with respect to diversification. Even when catastrophe bonds are cheaper than

[11] An interesting feature of the model is that contracts under homogeneity will always promise (but sometimes fail to deliver) full indemnification if the loss amount is less than company assets. This contrasts with the insurance demand model of Doherty and Schlesinger [8], in which the consumer buys less than full coverage in the presence of default risk. There are several reasons for this difference, including the presence of frictional costs in our model (Doherty and Schlesinger's result referenced actuarially fair premiums), and also a different definition of default: Doherty and Schlesinger's default features zero payment, while our model offers a partial pro-rata payment to insureds. The pro rata allocation amounts to a "mutuality principle" that encourages full coverage even in the presence of default risk.

insurance company assets (i.e., if $\delta_B < \delta_A$), they could still be strictly suboptimal if the welfare-maximizing solution involved tolerance of default beyond the absolute worst case scenario of N losses.

More generally, the return to catastrophe bonds given in (12) can be rewritten for the case of different frictional costs as:

$$\boxed{R_i = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x})\,[1 - f_{\mathbf{x}}] \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \sum_{j\in\Gamma(\mathbf{x})} w_j^{\mathbf{x}} \frac{\partial U_j^{\mathbf{x}}}{\partial W} \right) - \sum_{\mathbf{x}\notin\Omega^i} \lambda_{\mathbf{x}} + \mu(\delta_A - \delta_B) \le 0} \quad (13)$$

If bond principal enjoys frictional cost advantages relative to insurer assets ($\delta_B < \delta_A$), additional opportunities for catastrophe bond issuance may arise. Note, however, that such frictional cost advantages (to the extent they exist) are by no means the only consideration in assessing the potential for the catastrophe bond market. The extent of heterogeneity across consumers—in terms of preferences and in terms of risk exposures—as well as the presence or absence of diversification opportunities, still figure in the calculus of catastrophe bond issuance.

Catastrophe bonds are often advanced as a method for sidestepping the frictions in the reinsurance market. But, of course, a different set of frictional costs exist in the catastrophe bond market. The key policy question concerns whether supply-side initiatives to promote catastrophe risk transfer are best focused on the frictional costs in the reinsurance market or those in the catastrophe bond market.

At first glance, it seems that the opportunities for welfare gains are much greater when reducing frictional costs in the reinsurance market. The value of reducing the frictional costs of insurer assets by one unit is the derivative of the Lagrangian of the welfare function with respect to δ_A, or,

$$V_A \equiv \mu A^*,$$

while an analogous reduction in the frictional costs of catastrophe bond principal yields

$$V_B \equiv \mu \sum_i B_i^*.$$

Thus, the marginal benefit of reducing frictional costs in each market is directly proportional to the assets deployed in the respective market: Since far more collateral is held in the form of reinsurer assets than in the form of catastrophe bond principal, the marginal impact of frictional cost reductions in the reinsurance market should be far greater than similar reductions in the catastrophe bond market.

On the other hand, the cost side of the policy equation—i.e., what resources must be sacrificed to reduce frictional costs in each market—is less clear. Indeed, the frictional cost reduction technologies could differ substantially across the markets. Since the catastrophe bond market is young, there may be "low-hanging fruit": Investments in investor education or primary and secondary bond market infrastructure could offer much larger frictional cost reductions in the catastrophe bond market than could be possible in the more mature reinsurance market. However, frictional cost reductions will not necessarily translate perfectly into corresponding movements in market performance. The drawback of diversification inefficiencies discussed above may place important limits on the extent to which catastrophe risk will be transferred through catastrophe bonds and similar vehicles.

4.2 Optimal Collateralization with Generalized Capital Costs

To this point, we have abstracted from modelling any risk premium that might be demanded by investors. However, the amount of compensation due to capital suppliers theoretically depends on how the insurance risks borne correlate with returns on other capital assets. Little evidence exists connecting the cost of capital in the insurance industry to the risk characteristics of the underlying policyholder liabilities (Cummins and Harrington [5]; Cox and Rudd [4]); this is true of catastrophe insurance in particular, where both academics (e.g., Hoyt and McCullough [14]) and practitioners (e.g., Litzenberger, Beaglehole, and Reynolds [16]) have found catastrophic risk to be uncorrelated with stock market returns. Nevertheless, a connection exists at least in theory, so we explore this possibility more thoroughly in this subsection—where the cost of risk transfer depends not only on frictional costs associated with collateral but also on the relationship between the risk transferred and the risks in the broader capital markets.

In principle, frictional costs could be zero, and the cost of risk transfer could thus consist entirely of the equilibrium compensation for risk demanded by investors. However, as noted earlier, the cost of risk transfer must include a pure frictional cost that is entirely lost to the capital suppliers for the optimal collateralization structure to be determinate. If there is no gain to economizing on collateral, it does not matter how risk transfer is collateralized: All consumers will be fully insured, and it does not matter how insurance policies and catastrophe bonds are combined in providing that full insurance. However, if frictional costs are present, the earlier results still carry through—even after the introduction of security markets and equilibrium risk pricing.

To add security markets and equilibrium to our model, we start with a state-pricing approach built on the assumption of no arbitrage, as described in the first chapter of Duffie [9]. For insurance markets to be relevant (and not redundant), it must be the case that financial markets are incomplete. Specifically, if it is possible to replicate the payoffs from an insurance policy using existing securities in frictionless markets (and if those securities are priced fairly), there will be no need for insurance policies. Therefore, we will work under the assumption that markets are incomplete—or, specifically, that insurance policy payoffs cannot be replicated using other financial instruments. This is similar in flavor to the assumption underlying Mayers and Smith [18]—tradeable portfolio securities are assumed to have well-defined prices that flow from an equilibrium asset pricing model, but insurance risks are denoted as "nontradeable" and thus must be dealt with separately (although, as noted by Mayers and Smith, the decisions are not independent).

A key question not addressed by Mayers and Smith, but crucial for our purposes, is how insurance policies will be priced when they are non-tradeable and cannot be replicated with other financial instruments. The obvious approach is to price insurance policies "as if" they were traded financial securities. That is, their prices are determined by their contingent payoffs, weighted by appropriate state prices—just as with any other security. The only complication is that, by assumption, the contingencies relevant for insurance policy payoffs do not map into the state space that governs the security markets—so the state prices needed to price policies will not follow from the absence of arbitrage among the financial instruments traded in the security markets. With this in mind, we "extend" state prices derived from the assumption of no-arbitrage in the security markets to apply to subsets of events within states. While this extension may not be a technical implication of arbitrage pricing theory, it produces a necessary foundation for insurance pricing that is logically consistent with security market pricing.

We use the N person model and start by defining the relevant probability spaces. Recall the set of vectors that define consumer loss experience, denoted by Ω, whose members are row vectors $\mathbf{x}$ of length N, with the vector elements all taking a value of zero or one: $\mathbf{x}(i) = 1$ means that consumer i experienced a loss, while $\mathbf{x}(i) = 0$ means that she did not.

We now introduce a set of M securities, each security with distinct payoffs in S states of the world. Let Ψ be the set of those states (with the associated σ-algebra F_Ψ, and let there be state prices, consistent with the absence of arbitrage, denoted by π_s for each $s \in \Psi$. Let D be an $M \times S$ matrix, with D_{ij} describing the payoff of the i-th security in the j-th state. We assume that

$$span(D) \equiv \mathbb{R}^S.$$

This condition is typically known as a "complete markets" condition—that any arbitrary menu of state-contingent consumption can be purchased at time zero. In our case, however, it would be misleading to characterize markets as complete, since Ψ does not provide a complete description of the states of the world.

Instead, we characterize the full probability space as $(\Theta, F_\Theta, \mu_\Theta)$, with

$$\Theta \equiv \{\boldsymbol{\theta} = [\mathbf{x}(1)\ \mathbf{x}(2)...\mathbf{x}(N)\ s] \mid \mathbf{x} \in \Omega, s \in \Psi\}$$

The state variable $\boldsymbol{\theta} \in \Theta$ is a row vector of length $N + 1$ that provides a complete description of one possible state of the world. The first N elements of $\boldsymbol{\theta}$ describe which consumers experienced losses (and which ones did not), while the last element describes the state of the securities markets. The entire set Θ contains all possible states of the world.

The following set definitions are useful:

$$\Theta^i = \{\boldsymbol{\theta} : \boldsymbol{\theta}(i) = 1\},$$

the set of all states in which agent i suffers a loss, and

$$\Gamma(\boldsymbol{\theta}) = \{i : \boldsymbol{\theta}(i) = 1\},$$

the set of all agents that lose in state $\boldsymbol{\theta}$. In addition, for every s and every agent i, define:

$$\Upsilon_i^{s0} \equiv \{\boldsymbol{\theta} : \boldsymbol{\theta} \notin \Theta^i, \boldsymbol{\theta}(N+1) = s\},$$

the set of all states $\boldsymbol{\theta}$ where agent i does not suffer a loss and the security market "sub-state" is s, and

$$\Upsilon_i^{s1} \equiv \{\boldsymbol{\theta} : \boldsymbol{\theta} \in \Theta^i, \boldsymbol{\theta}(N+1) = s\},$$

the set of all states $\boldsymbol{\theta}$ where agent i *does* suffer a loss and the security market "sub-state" is s. Finally, note that for any i the entire sub-state space defined by s can be written as:

$$\Upsilon^s \equiv \Upsilon_i^{s1} \cup \Upsilon_i^{s0}$$

We now extend the state prices to define prices for events that are not measurable with respect to F_Ψ. Define extended state prices as follows. For each $s \in \Psi$

$$\pi^\theta = \pi_s, \forall \boldsymbol{\theta} \in \Theta \mid \boldsymbol{\theta}(N+1) = s$$

Recall that we are extending the state-prices to include the price of claims associated with the hazards being insured. This approach implicitly assumes that there is no variation in "sub-state" prices within the states priced by the security market equilibrium. This is not a mathematical truism implied by the absence of arbitrage in security markets: The absence of arbitrage does not pin down the state prices for events that are not measurable with respect to F_Ψ. However, as discussed above, this assumption provides a reasonable basis for insurance pricing that is logically consistent with the security market equilibrium.

We can now define utility for Consumer i as

$$EU_i = \sum_{s\in\Psi} \sum_{\boldsymbol{\theta}\in\Upsilon_i^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_i \left(W_{is} - L_i + f_{\boldsymbol{\theta}} I_i + B_i\right) + \sum_{s\in\Psi} \sum_{\boldsymbol{\theta}\in\Upsilon_i^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_i \left(W_{is}\right),$$

where $f_{\boldsymbol{\theta}}$ represents the proportion of the indemnity payment actually paid in state $\boldsymbol{\theta}$, and W_{is} is state-contingent security-market wealth for consumer i.

The social planning problem can now be written as:

$$\max_{A,\{B_i\},\{c_i\},\{I_i\},\{f_{\boldsymbol{\theta}}\},\{W_{is}\}} V = \sum_i EU_i \tag{14}$$

subject to

$$[\mu] \quad : \quad \sum c_i \geq \delta_A A + \delta_B \sum_i B_i + \sum_{\boldsymbol{\theta}\in\Theta} \left(\mu_\Theta(\boldsymbol{\theta}) \pi^\theta \sum_{i\in\Gamma(\boldsymbol{\theta})} (f_{\boldsymbol{\theta}} I_i + B_i) \right) \tag{15}$$

$$[\lambda_{\boldsymbol{\theta}}] \quad : \quad f_{\boldsymbol{\theta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} I_i \leq A, \forall \boldsymbol{\theta} \tag{16}$$

$$[\phi_{\boldsymbol{\theta}}] \quad : \quad f_{\boldsymbol{\theta}} \leq 1, \forall \boldsymbol{\theta} \tag{17}$$

$$[\varphi_i] \quad : \quad \sum_{s\in\Psi} \pi_s W_{is} \leq W - c_i, \forall i \tag{18}$$

and the previous non-negativity constraints. New to the problem is the last constraint, which governs portfolio investment. The present value of each consumer's contingent consumption is constrained to no more than time zero wealth, net of risk transfer costs.

Before proceeding further, it is worth noting that the generalized capital allocation structure does not alter our basic results for the return on catastrophe bonds (see Appendix Section C.1). In this setting, the return to catastrophe bonds is given by:

$$\boxed{R_i = \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\theta)(1-f_\theta) \left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \sum_{i\in\Gamma(\theta)} w_i^\theta \frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} \right) - \sum_{\theta\notin\Theta^i} \lambda_\theta + \mu(\delta_A - \delta_B) \leq 0} \tag{19}$$

This is analogous to equation 13.

The marginal conditions describing optimal issuance of catastrophe bonds do not directly reference frictional costs (except when $\delta_B \neq \delta_A$), yet frictional costs are required for them to be useful. Without frictional costs, both terms of the marginal condition will be zero in all cases—simply because there is no insurance company default in the absence of frictional costs, and every consumer enjoys full insurance.

To illustrate this, we eliminate frictional costs but continue to weight claims payments and catastrophe bond default according to the appropriate state prices:

$$\sum c_i \geq \sum_{\boldsymbol{\theta}\in\Theta} \left(\mu_\Theta(\boldsymbol{\theta})\pi^\theta f_{\boldsymbol{\theta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} I_i \right) + \sum_i \left(\sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \right) B_i.$$

Intuitively, insurance premiums must pay for contingent claims on the assets, but, in the absence of frictional costs, extra capital always earns its own keep, because the insurer can invest it at market rates of return. As a result, the only real cost of holding capital as collateral is the extent to which you are exposing the owner to additional claims risk—but this additional risk is compensated in the capital market.

Section C.2 of the Appendix characterizes the solution in this setting without frictional costs in a series of lemmas that yield two important lessons. First, full insurance will always be optimal: Every consumer will be fully covered. Second, the optimal division of collateral between reinsurance company assets and catastrophe bond principal is indeterminate. In other words, with full insurance it ends up being irrelevant whether the insurance is provided through reinsurance policies or catastrophe bonds: Any combination of the two instruments is optimal, so long as full insurance is provided.

As is often the case, consumers fully insure when the insurance is fairly priced. Moreover, there is no incentive to economize on collateral in this setting—collateral is "free" in the sense that there are no frictional costs. The only "costs" with holding assets in the reinsurer or the SPV are the fair value of expected claims payments to policyholders or cat bond issuers. Since there are no penalties associated with over-collateralization, neither risk transfer instrument holds a natural advantage over the other.

5 Other Risk Transfer Options

To this point, we have limited our attention to catastrophe bonds and traditional reinsurance policies. With this focus, we risk overlooking hedging strategies based on other risk transfer options that could potentially yield welfare improvements. In this section, we consider how other risk transfer strategies fit into our framework.

5.1 Collateralized Reinsurance and Side-cars

Reinsurance can be "collateralized" in at least two senses. The first, more common sense, is a contract clause requiring the reinsurer to collateralize claims obligations at the time they are incurred but before they are due to be paid. The second is a full or partial collateralization of the policy limits at the inception of the contract. Both are interesting for our paper because they allow the reinsurer and its customers to influence how assets are divided up in the event of bankruptcy. We discuss each in turn.

Reinsurance contract clauses regarding the collateralization of liabilities arise most often in the context of transactions between offshore reinsurers and U.S. cedents (i.e., buyers of reinsurance). Regulations regarding statutory credit for reinsurance typically stipulate that an insurer may not take credit for anticipated recoveries from unlicensed reinsurers unless those anticipated recoveries are fully secured. Acceptable forms of security include funds held in trust and clean, irrevocable, and evergreen letters of credit issued by financial institutions deemed acceptable by the regulator in question.

To the extent that some cedents have these contract clauses and others do not, the clauses may be interpreted as a means of affecting the distribution of assets in bankruptcy. Secured claimants effectively "step ahead" of unsecured claimants in the liquidation process. However, it should be noted that the ability to "step ahead" is by no means absolute and depends on ex post actions by the insurer. For example, a transfer of assets to a trust for the benefit of a cedent (or to collateralize a letter of credit issued by a third-party for the benefit of a cedent) can be challenged as a voidable preference if bankruptcy follows soon thereafter.[12] Hence, in practice, a cedent cannot count on security being posted when a reinsurer is in or near insolvency, and, even if the reinsurer is willing to post security—the transfer will be subject to challenge. This issue also surfaces when modelling collateralization in a one-period setting: When claims are submitted simultaneously and bankrupt the reinsurer, what compels the reinsurer (or its receiver) to honor liability collateralization clauses that have no force during liquidation?

It is also possible to provide partial or full collateralization (e.g., a letter of credit) of the policy limits at contract inception. This approach is useful if the underwriter does not have a financial strength rating, as is the case with a number of hedge funds that have entered the property catastrophe reinsurance market in recent years. For an underwriter issuing only fully collateralized policies, this form of collateralized reinsurance is effectively equivalent to a catastrophe bond with an indemnity trigger. Another variation on this theme is the recent innovation of the reinsurance "side-car," where investors capitalize a special purpose company to provide reinsurance to a ceding company, with the capital being held in a collateral trust account for the benefit of the cedent (see Murray [19] for additional information on side-cars).

For underwriters issuing policies that differ in the degree of collateralization, the situation is more complicated. Since collateral posted will presumably be released in the event that the underlying policy is not triggered, it will subsequently become available to pay claimants whose policy limits were not fully secured at inception. In our framework, this approach to collateralizing risk transfer offers the potential for welfare improvement relative to catastrophe bonds because of this increase in the availability of assets to pay claims.

In principle, varying the degree of collateralization across policies could be used to affect the allocation of assets during bankruptcy, but it is important to note theoretical limits. Aside from the two-person case, varying the degree of collateralization associated with policies will generally give the insurer only limited control over the allocation of assets in

[12]For more details, especially with respect to letters of credit, see Hall [13] and the NAIC's *Receivers Handbook for Insurance Company Insolvencies.*

bankruptcy. In the multiple consumer case, there are multiple bankruptcy states, with different consumers affected in different bankruptcy states: Varying ex ante collateral levels does not allow arbitrary prioritization of claims within bankruptcy.

Practical limits also apply. The effective security of partial collateralization will not be transparent to the policyholder in practice: To form an expectation of relative priority in bankruptcy, one must know all the details about the collateralization of other policies. For example, if all outstanding policies are collateralized to the same degree, collateralization has no effect, and recoveries would not change if the collateralization were terminated. Finally, in a world where claims are being submitted in continuous time, the reinsurer will not be able to commit all of its assets to ex ante collateralization, since it would have no funds available beyond the collateral supporting any given policy to pay a claim on that policy.

5.2 Third-Party Default Insurance

In the model, catastrophe bonds protect consumers from the consequences of insurer default. A possible alternative approach would be to buy "default insurance" from a third party, and these possibilities arise if the insurer has issued negotiable securities. For example, credit default swaps (CDS) referencing outstanding transferable bonds or loans of the insurer[13] are a possible alternative hedging device: Issuing policies along with default protection could conceivably offer welfare improvements over the strategies considered in this paper. There are two aspects of this strategy that merit comment.

First, to deliver a welfare improvement over catastrophe bonds, any protection offered by the CDS would have to be implicitly collateralized at less than 100% (or somehow have lower frictional costs with a similar degree of collateralization). On the one hand, less than full collateralization implies that there is some risk of counterparty default, so purchasing default insurance with a CDS is not a perfect substitute for the catastrophe bond. On the other hand, if the seller of protection could conceivably realize additional economies in collateralization by taking advantage of diversification opportunities beyond the insurance market. In this light, the CDS approach can be understood as an intermediate device, located on a continuum between the fully collateralized and undiversified catastrophe bond, and the imperfectly collateralized and diversified reinsurance contract.

Second, the CDS hedging approach will generally involve basis risk when there are more than two policyholders. With two policyholders, the CDS offers a perfect hedge to a consumer desiring protection: The contract is triggered only in states where the company defaults *and* the consumer experiences a loss. With multiple consumers, however, the company may default in states where the consumer does not experience a loss, thereby triggering a default insurance payment in a scenario where the consumer does not need additional indemnification. The consumer becomes "overinsured" with respect to default risk. Moreover, the extent of a consumer's recovery on a policy will generally vary across states of default in the multiple consumer model, and this variation may not generally be hedged with a CDS (or with a short position in the insurer's debt, if this were possible) *unless* the holder of the underlying debt security is in the same class as the holder of an

[13]Alternatively, in the absence of a CDS market, hedging strategies using equities or equity derivatives could substitute. However, equity-based strategies will generally have more basis risk as defined below.

insurance policy with respect to priority of claim on the insurer's assets.[14]

6 Conclusions

In theory, catastrophe bonds can potentially be useful to ameliorate the effects of the contracting constraints faced by insurers. These constraints include the difficulty of writing contracts that allocate the company's assets efficiently in the event of insolvency, or of contracts that are contingent on the loss experiences of other insureds. These constraints can bind when insureds are heterogeneous. Therefore, catastrophe bonds can be welfare-improving when: (1) Reinsurers face constraints on contracting, and (2) Insureds are heterogeneous. We have derived these results from models of efficient collateral allocation with two or more insureds, when the frictional costs associated with catastrophe bond issuance mirror those associated with holding assets in insurance companies.

If catastrophe bond issuance is a cheaper option than equity issuance (with respect to frictional costs), additional opportunities for welfare-improvement arise. Indeed, if bond issuance costs (i.e., spreads over LIBOR that must be paid by issuers as premium) fall, it seems likely that usage will expand. However, because of the catastrophe bond's relative inefficiency in the realm of diversification, it is possible for the catastrophe bond to be cheaper but still unattractive relative to risk transfer options based on partial collateralization. Thus, while frictional costs associated with underdeveloped market infrastructure and the basis risk faced by issuers are often fingered as the main roadblocks to growth in the catastrophe bond market, this analysis suggests that a more fundamental obstacle—costs deriving from the instrument's full collateralization—may ultimately place limits on its growth in the absence of further innovation.

[14]This equivalence in priority will often hold for reinsurance contracts, which are in the same class as general creditors under the NAIC's Insurer Rehabilitation and Liquidation Model Act (revision of 4/27/04) but not for primary insurance contracts, which are typically assigned higher priority. In some circumstances, however, reinsurance contracts have also been assigned higher priority than general creditors (see Hall [12]).

APPENDIX

A Unconstrained Contracting

In this section of the Appendix, we verify that catastrophe bond issuance is optimally zero when insurance policy indemnity schedules can be made state-contingent and frictional costs for the two risk transfer technologies are equal.

Let $I_i^{\mathbf{x}}$ represent the indemnity promised in state $\mathbf{x}$ under the policy issued to consumer i. Then the utility of consumer i can be rewritten as:

$$EU_i = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) U_i \left(W_i - L_i + f_{\mathbf{x}} I_i^{\mathbf{x}} + B_i - c_i\right) + \sum_{\mathbf{x}\notin\Omega^i} \Pr(\mathbf{x}) U_i \left(W_i - c_i\right), \qquad (20)$$

and the social planning problem becomes:

$$\max_{A,\{B_i\},\{c_i\},\{I_i^{\mathbf{x}}\},\{f_{\mathbf{x}}\}} V = \sum_i EU_i \qquad (21)$$

subject to:

$$[\mu] \;:\; \sum c_i \geq \delta_A A + \delta_B \sum_i B_i + \sum_{\mathbf{x}\in\Omega} \Pr(\mathbf{x}) \left(\sum_{i\in\Gamma(\mathbf{x})} f_{\mathbf{x}} I_i^{\mathbf{x}} + \sum_{i\in\Gamma(\mathbf{x})} B_i \right) \qquad (22)$$

$$[\lambda_{\mathbf{x}}] \;:\; f_{\mathbf{x}} \sum_{i\in\Gamma(\mathbf{x})} I_i^{\mathbf{x}} \leq A, \forall \mathbf{x} \qquad (23)$$

$$[\phi_{\mathbf{x}}] \;:\; 0 \leq f_{\mathbf{x}} \leq 1, \forall \mathbf{x} \qquad (24)$$

and the usual non-negativity constraints and subject to $I_i^{\mathbf{x}} = 0$ for $\mathbf{x}\notin\Omega^i$, for all i. The following theorem proves the result.

Theorem. *Suppose $\delta_A = \delta_B = \delta$. Suppose $\hat{A}, \{\hat{B}_i\}, \{\hat{c}_i\}, \{\hat{I}_i^{\mathbf{x}}\}, \{\hat{f}_{\mathbf{x}}\}$ is a set of optimal choices maximizing (21) subject to the listed constraints, with $\hat{B}_k > 0$ for at least one $k \in \{1, ..., N\}$. Then there exists another set of choices $\bar{A}, \{\bar{B}_i\}, \{\bar{c}_i\}, \{\bar{I}_i^{\mathbf{x}}\}, \{\bar{f}_{\mathbf{x}}\}$, with $\bar{B}_i = 0$ for all i, that satisfy the listed constraints and also maximize the objective function.*

Proof:

Define $\Lambda = \{k : \text{k} \in \{1, ..., N\}, \hat{B}_k > 0\}$.

Let $\bar{c}_i = \hat{c}_i$, for all i.

Let $\bar{B}_i = 0$, for all i.

Let $\bar{f}_{\mathbf{x}} = 1$, for all $\mathbf{x}$.

Let $\bar{A} = \hat{A} + \sum_i \hat{B}_i$

For all i : Let $\bar{I}_i^{\mathbf{x}} = \hat{f}_{\mathbf{x}} \hat{I}_i^{\mathbf{x}} + \hat{B}_i$, $\forall \mathbf{x} \in \Omega^i$; $\bar{I}_i^{\mathbf{x}} = \hat{I}_i^{\mathbf{x}} = 0$, $\forall \mathbf{x} \notin \Omega^i$.

We first verify that the new choices satisfy all constraints:

Note that

$$\begin{aligned}\sum \bar{c}_i &= \sum \hat{c}_i \geq \delta\left(\bar{A}+\sum_i \bar{B}_i\right)+\sum_{\mathbf{x}\in\Omega}\Pr(\mathbf{x})\left(\sum_{i\in\Gamma(\mathbf{x})}\bar{f}_{\mathbf{x}}\bar{I}_i^{\mathbf{x}}+\sum_{i\in\Gamma(\mathbf{x})}\bar{B}_i\right)\\ &= \delta\left(\hat{A}+\sum_i \hat{B}_i\right)+\sum_{\mathbf{x}\in\Omega}\Pr(\mathbf{x})\left(\sum_{i\in\Gamma(\mathbf{x})}\left(\hat{f}_{\mathbf{x}}\hat{I}_i^{\mathbf{x}}+\hat{B}_i\right)\right)\\ &= \delta\left(\hat{A}+\sum_i \hat{B}_i\right)+\sum_{\mathbf{x}\in\Omega}\Pr(\mathbf{x})\left(\sum_{i\in\Gamma(\mathbf{x})}\hat{f}_{\mathbf{x}}\hat{I}_i^{\mathbf{x}}+\sum_{i\in\Gamma(\mathbf{x})}\hat{B}_i\right),\end{aligned}$$

so $[\mu]$ is satisfied.

For all $\mathbf{x}$,

$$\bar{f}_{\mathbf{x}}\sum_{i\in\Gamma(\mathbf{x})}\bar{I}_i^{\mathbf{x}}=\sum_{i\in\Gamma(\mathbf{x})}\left(\hat{f}_{\mathbf{x}}\hat{I}_i^{\mathbf{x}}+\hat{B}_i\right)$$

But

$$\bar{A}=\hat{A}+\sum_i \hat{B}_i \geq \sum_{i\in\Gamma(\mathbf{x})}\left(\hat{f}_{\mathbf{x}}\hat{I}_i^{\mathbf{x}}+\hat{B}_i\right)$$

since

$$\hat{A}\geq \sum_{i\in\Gamma(\mathbf{x})}\hat{f}_{\mathbf{x}}\hat{I}_i^{\mathbf{x}},$$

so $[\lambda_{\mathbf{x}}]$ is satisfied.

Since $\bar{f}_{\mathbf{x}}=1$, for all $\mathbf{x}$, $[\phi_{\mathbf{x}}]$ is satisfied, and the non-negativity constraints are obviously satisfied since the original set of choices met the non-negativity constraints. So the proposed choices are feasible ones.

Now we show that the proposed alternative choices produce the same value of the objective function. Note that

$$EU_i=\sum_{\mathbf{x}\in\Omega^i}\Pr(\mathbf{x})U_i\left(W_i-L_i+\bar{f}_{\mathbf{x}}\bar{I}_i^{\mathbf{x}}+\bar{B}_i-\bar{c}_i\right)+\sum_{\mathbf{x}\notin\Omega^i}\Pr(\mathbf{x})U_i\left(W_i-\bar{c}_i\right)$$

can be rewritten as:

$$\sum_{\mathbf{x}\in\Omega^i}\Pr(\mathbf{x})U_i\left(W_i-L_i+\hat{f}_{\mathbf{x}}\hat{I}_i^{\mathbf{x}}+\hat{B}_i-\hat{c}_i\right)+\sum_{\mathbf{x}\notin\Omega^i}\Pr(\mathbf{x})U_i\left(W_i-\hat{c}_i\right),$$

so, for each consumer, the proposed alternative choices yield the exact same utility as the original choices. Therefore, the objective function value does not change when we move from the original choices to the new choices. Q.E.D.

B Proofs for Section 4.1

B.1 Derivation of (12) and (13)

We start by developing the first order conditions from the stated maximization problem. We then use these to derive Equation 13 (which reduces to Equation 12 when $\delta_A = \delta_B$). The first order conditions are as follows (where we use the notation $U_i^{\mathbf{x}}$ to denote the utility of consumer i in state $\mathbf{x}$):

$$[B_i] \quad : \quad \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) \frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \left(\delta_B + \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) \right) \leq 0 \tag{25}$$

$$[c_i] \quad : \quad -\sum_{\mathbf{x}\in\Omega} \Pr(\mathbf{x}) \frac{\partial U_i^{\mathbf{x}}}{\partial W} + \mu = 0 \tag{26}$$

$$[I_i] \quad : \quad \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) f_{\mathbf{x}} \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \right) - \sum_{\mathbf{x}\in\Omega^i} f_{\mathbf{x}} \lambda_{\mathbf{x}} = 0 \tag{27}$$

$$[A] \quad : \quad -\delta_A \mu + \sum_{\mathbf{x}\in\Omega} \lambda_{\mathbf{x}} = 0 \tag{28}$$

$$[f_{\mathbf{x}}] \quad : \quad \sum_{i\in\Gamma(\mathbf{x})} \Pr(\mathbf{x}) I_i \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \right) - \lambda_{\mathbf{x}} \sum_{i\in\Gamma(\mathbf{x})} I_i - \phi_{\mathbf{x}} = 0 \tag{29}$$

We start by observing that $\phi_{\mathbf{x}} = 0$ for all states $\mathbf{x}$ (i.e., $\forall \mathbf{x}$, the constraint $f_{\mathbf{x}} \leq 1$ fails to bind). To see this, multiply $[I_i]$ by I_i and sum over i to obtain:

$$\sum_{i=1}^{N} \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) f_{\mathbf{x}} I_i \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \right) - \sum_{i=1}^{N} \sum_{\mathbf{x}\in\Omega^i} f_{\mathbf{x}} I_i \lambda_{\mathbf{x}} = 0 \tag{30}$$

Next, multiply $[f_{\mathbf{x}}]$ by $f_{\mathbf{x}}$ and sum over $\mathbf{x}$ to obtain:

$$\sum_{\mathbf{x}\in\Omega} \sum_{i\in\Gamma(\mathbf{x})} \Pr(\mathbf{x}) f_{\mathbf{x}} I_i \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \right) - \sum_{\mathbf{x}\in\Omega} \sum_{i\in\Gamma(\mathbf{x})} f_{\mathbf{x}} I_i \lambda_{\mathbf{x}} - \sum_{\mathbf{x}\in\Omega} f_{\mathbf{x}} \phi_{\mathbf{x}} = 0 \tag{31}$$

After noting that $\Gamma(\mathbf{x})$ is a null set for the state where nobody experiences a loss (i.e., where $\mathbf{x}$ is a vector of zeroes), it is evident that the first two terms of (30) are equal to the first two terms of (31),[15] implying that $\sum_{\mathbf{x}\in\Omega} f_{\mathbf{x}} \phi_{\mathbf{x}} = 0$. Thus, it is clear that $\phi_{\mathbf{x}} = 0$ for all $\mathbf{x}$.

We now derive the marginal utility of catastrophe bond issuance for the case of two consumers. After multiplying by $f_{\mathbf{x}}$, using the above result on $\phi_{\mathbf{x}}$, and rearranging, note that $[f_{\mathbf{x}}]$ can be written as:

$$\sum_{i\in\Gamma(\mathbf{x})} \Pr(\mathbf{x}) w_i^{\mathbf{x}} f_{\mathbf{x}} \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \right) - f_{\mathbf{x}} \lambda_{\mathbf{x}} = 0,$$

[15] Intuitively, the summations $\sum_{i=1}^{N} \sum_{\mathbf{x}\in\Omega^i}$ represent the sum of all states in which agent i suffers a loss, summed across all agents i. This is equivalent to the sum of all agents suffering a loss in state $\mathbf{x}$, summed across all states $\mathbf{x}$, which is depicted by the double summation $\sum_{\mathbf{x}\in\Omega} \sum_{i\in\Gamma(\mathbf{x})}$.

where

$$w_i^{\mathbf{x}} = \frac{I_i}{\sum_{i\in\Gamma(\mathbf{x})} I_i}.$$

Summing over $\mathbf{x} \in \Omega^i$:

$$\sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) f_{\mathbf{x}} \left(\sum_{j\in\Gamma(\mathbf{x})} w_j^{\mathbf{x}} \frac{\partial U_j^{\mathbf{x}}}{\partial W} - \mu \right) - \sum_{\mathbf{x}\in\Omega^i} f_{\mathbf{x}} \lambda_{\mathbf{x}} - \sum_{\mathbf{x}\in\Omega^i} \frac{f_{\mathbf{x}} \phi_{\mathbf{x}}}{\sum_{j\in\Gamma(\mathbf{x})} I_j} = 0,$$

which is the same as $[I_i]$ except that $\frac{\partial U_i^{\mathbf{x}}}{\partial W}$ in each state is replaced by $\sum_{i\in\Gamma(\mathbf{x})} w_i^{\mathbf{x}} \frac{\partial U_i^{\mathbf{x}}}{\partial W}$ (a weighted average of the marginal utilities of all consumers who lost in that state). In summary, we have

$$\sum_{\mathbf{x}\in\Omega^i} f_{\mathbf{x}} \lambda_{\mathbf{x}} = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) f_{\mathbf{x}} \left(\sum_{j\in\Gamma(\mathbf{x})} w_j^{\mathbf{x}} \frac{\partial U_j^{\mathbf{x}}}{\partial W} - \mu \right). \tag{32}$$

Note that we did not need to multiply by $f_{\mathbf{x}}$. Omitting this step leads to:

$$\sum_{\mathbf{x}\in\Omega^i} \lambda_{\mathbf{x}} = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) \left(\sum_{j\in\Gamma(\mathbf{x})} w_j^{\mathbf{x}} \frac{\partial U_j^{\mathbf{x}}}{\partial W} - \mu \right) \tag{33}$$

Recall the marginal condition $[B_i]$:

$$R_i = -(\delta + \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}))\mu + \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) \frac{\partial U_i^{\mathbf{x}}}{\partial W} \leq 0.$$

The first term is the marginal cost of issuance—including both the frictional cost per dollar of collateral and the expected loss on the bond—and the second term is the marginal benefit, which amounts to an extra dollar of consumption in all of the loss states. Subtracting the left-hand side of the first order condition $[I_i]$ from the above expression yields the following:

$$R_i = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) \left[1 - f_{\mathbf{x}}\right] \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \right) + \sum_{\mathbf{x}\in\Omega^i} f_{\mathbf{x}} \lambda_{\mathbf{x}} - \delta_B \mu.$$

or

$$R_i = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) \left[1 - f_{\mathbf{x}}\right] \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \right) + \sum_{\mathbf{x}\in\Omega^i} \left[f_{\mathbf{x}} - 1\right] \lambda_{\mathbf{x}} + \sum_{\mathbf{x}\in\Omega^i} \lambda_{\mathbf{x}} - \delta_B \mu.$$

Substituting in from $[A]$ yields:

$$R_i = \sum_{\mathbf{x}\in\Omega^i} \Pr(\mathbf{x}) \left[1 - f_{\mathbf{x}}\right] \left(\frac{\partial U_i^{\mathbf{x}}}{\partial W} - \mu \right) + \sum_{\mathbf{x}\in\Omega^i} \left[f_{\mathbf{x}} - 1\right] \lambda_{\mathbf{x}} - \sum_{\mathbf{x}\notin\Omega^i} \lambda_{\mathbf{x}} + \left(\delta_A - \delta_B\right) \mu.$$

Subtracting (33) from (32) and substituting in the resulting expression for $\sum_{\mathbf{x}\in\Omega^i} \left[f_{\mathbf{x}} - 1\right] \lambda_{\mathbf{x}}$ yields Equation 13.

B.2 Irrelevance of Cat Bonds in the Absence of Default Risk

In Footnote 10, we claim that a solution with positive cat bond issuance for consumer i, where that consumer faces no default risk, is not unique and can be replicated by a solution with zero issuance for consumer i. The following theorem proves this claim formally.

Theorem. *Suppose $\delta_A = \delta_B = \delta$ and that $\hat{A}, \{\hat{B}_i\}, \{\hat{c}_i\}, \{\hat{I}_i\}, \{\hat{f}_{\mathbf{x}}\}$ is a set of optimal choices maximizing (21) subject to the listed constraints, with $\hat{B}_k > 0$ for some consumer $k \in \{1, ..., N\}$, and $\hat{f}_{\mathbf{x}} = 1$ for all $\mathbf{x} \in \Omega^k$. Then there exists another set of choices $\bar{A}, \{\bar{B}_i\}, \{\bar{c}_i\}, \{\bar{I}_i\}, \{\bar{f}_{\mathbf{x}}\}$, with $\bar{B}_k = 0$ that satisfy the listed constraints and also maximize the objective function.*

Proof:

Let $\bar{c}_i = \hat{c}_i$, for all i.
Let $\bar{B}_i = \hat{B}_i$, for all $i \neq k$. Let $\bar{B}_k = 0$.
Let $\bar{I}_i = \hat{I}_i$, for all $i \neq k$. Let $\bar{I}_k = \hat{I}_k + \hat{B}_k$
Let $\bar{f}_{\mathbf{x}} = \hat{f}_{\mathbf{x}}$, for all $\mathbf{x}$.
Let $\bar{A} = \hat{A} + \hat{B}_k$
We first verify that all constraints are satisfied. Note that

$$\sum \bar{c}_i = \sum \hat{c}_i \geq \delta \left(\hat{A} + \sum_i \hat{B}_i \right) + \sum_{\mathbf{x} \in \Omega} \Pr(\mathbf{x}) \left(\sum_{i \in \Gamma(\mathbf{x})} \hat{f}_{\mathbf{x}} \hat{I}_i + \sum_{i \in \Gamma(\mathbf{x})} \hat{B}_i \right) =$$

$$\delta \left(\bar{A} + \sum_i \bar{B}_i \right) + \sum_{\mathbf{x} \in \Omega} \Pr(\mathbf{x}) \left(\sum_{i \in \Gamma(\mathbf{x})} \bar{f}_{\mathbf{x}} \bar{I}_i + \sum_{i \in \Gamma(\mathbf{x})} \bar{B}_i \right),$$

so $[\mu]$ is satisfied.

For all $\mathbf{x} \notin \Omega^k$, $[\lambda_{\mathbf{x}}]$ is obviously satisfied. For $\mathbf{x} \in \Omega^k$, we know that

$$\sum_{i \in \Gamma(\mathbf{x})} \bar{I}_i = \sum_{i \in \Gamma(\mathbf{x})} \hat{I}_i + \hat{B}_k,$$

and that $\bar{f}_{\mathbf{x}} = \hat{f}_{\mathbf{x}} = 1$ for all $\mathbf{x} \in \Omega^k$. So, it follows that

$$\bar{f}_{\mathbf{x}} \sum_{i \in \Gamma(\mathbf{x})} \bar{I}_i = \sum_{i \in \Gamma(\mathbf{x})} \hat{I}_i + \hat{B}_k \leq \hat{A} + \hat{B}_k = \bar{A},$$

so $[\lambda_{\mathbf{x}}]$ is also satisfied for $\mathbf{x} \in \Omega^k$.

Finally, the constraints $[\phi_{\mathbf{x}}]$ are obviously satisfied, since $\bar{f}_{\mathbf{x}} = \hat{f}_{\mathbf{x}}$, for all $\mathbf{x}$.

Next, we verify that the objective function value is unchanged with the new choices. Note that, for all i,

$$\sum_{\mathbf{x} \in \Omega^i} \Pr(\mathbf{x}) U_i \left(W_i - L_i + \bar{f}_{\mathbf{x}} \bar{I}_i + \bar{B}_i - \bar{c}_i \right) + \sum_{\mathbf{x} \notin \Omega^i} \Pr(\mathbf{x}) U_i \left(W_i - \bar{c}_i \right)$$

equals

$$\sum_{\mathbf{x} \in \Omega^i} \Pr(\mathbf{x}) U_i \left(W_i - L_i + \hat{f}_{\mathbf{x}} \hat{I}_i + \hat{B}_i - \hat{c}_i \right) + \sum_{\mathbf{x} \notin \Omega^i} \Pr(\mathbf{x}) U_i \left(W_i - \hat{c}_i \right).$$

Thus, expected utility for each consumer under the new choices is identical to that under the original choices. It follows that the overall objective function $V = \sum EU_i$ must have the same value under each set of choices. Q.E.D.

B.3 Homogeneity with Identical Frictional Costs

In this section, we assume that consumers are "homogenous" in the sense of having identical preferences and wealth, as well as having risks that are independent and identically distributed. In this homogeneous case, we restrict ourselves to considering symmetric equilibria in which all consumers are treated identically.

With this focus, the choice problem can be simplified by recognizing that all consumers have the same contracts and bond issuance. We denote the common probability of loss as p and the size of loss as L and seek to solve:

$$\max_{B,I,E,c,\{f_l\}} \sum_{l=0}^{N} \binom{N}{l} p^l (1-p)^{N-l} \left[\begin{array}{c} l * U(W - L + f_l I + (1-\delta-p)B - c) + \\ (N-l) * U(W - (\delta+p)B - c) \end{array} \right]$$

subject to the following constraints, with their associated multipliers,

$$\begin{aligned} [\mu] &: \; Nc \geq \delta A + \sum_{l=0}^{N} \binom{N}{l} p^l (1-p)^{N-l} l f_l I && (34) \\ [\lambda_l] &: \; l f_l I \leq A, \forall l && (35) \\ [\phi_l] &: \; f_l \leq 1, \forall l && (36) \end{aligned}$$

The contracting constraints are embodied in the f_l multipliers,which allow the indemnity payments to be scaled back in any state of the world (e.g., in states of default), but restricts any discounting to apply evenly across policyholders. Note that f_l is allowed to vary with l, the number of insureds experiencing a loss. Also, note that we have automatically allocated the frictional cost associated with cat bond issuance to the consumer (rather than including it in constraint $[\mu]$; this has no effect on the result.

The first order conditions for this problem are as follows (where we use the notation U_l' to denote the marginal utility of a consumer who experienced a loss along with $l-1$ other consumers, and $\bar{U}'$ is marginal utility in the no loss state):

$$[B] \quad : \quad \sum_{l=0}^{N} \binom{N}{l} p^l (1-p)^{N-l} \left[l U_l' (1-\delta-p) - (N-l)\bar{U}'(\delta+p) \right] \leq 0 \tag{37}$$

$$[c] \quad : \quad \sum_{l=0}^{N} \binom{N}{l} p^l (1-p)^{N-l} \left[-l U_l' - (N-l)\bar{U}' \right] + N\mu = 0 \tag{38}$$

$$[I] \quad : \quad \sum_{l=0}^{N} \binom{N}{l} p^l (1-p)^{N-l} \left(U_l' - \mu \right) l f_l - \sum l f_l \lambda_l \leq 0 \tag{39}$$

$$[A] \quad : \quad -\delta\mu + \sum \lambda_l \leq 0 \tag{40}$$

$$[f_l] \quad : \quad \binom{N}{l} p^l (1-p)^{N-l} \left(U_l' - \mu \right) l I - \lambda_l l I - \phi_l = 0 \tag{41}$$

Note that the optimality condition for $[c]$ can be used to rewrite $[B]$ as:

$$-(\delta+p)N\mu + \sum_{l=0}^{N} \binom{N}{l} p^l (1-p)^{N-l} l U_l' = -\delta\mu + \sum_{l=0}^{N} \binom{N}{l} p^l (1-p)^{N-l} \left(\frac{l}{N} \right) [U_l' - \mu] \, . \tag{42}$$

To rewrite this, we relied on the fact that, with N consumers,

$$p = p^1(1-p)^{N-1} + (N-1)p^2(1-p)^{N-2} + ... + p^N = \sum_{i=1}^{N} \binom{N-1}{i-1} p^i (1-p)^{N-i},$$

And, going further, that:

$$Np = N \sum_{i=1}^{N} \binom{N-1}{i-1} p^i (1-p)^{N-i} = N \sum_{i=1}^{N} \frac{N-1!}{N-i!i-1!} p^i (1-p)^{N-i} = \sum_{i=1}^{N} \binom{N}{i} i p^i (1-p)^{N-i}$$

But $[A]$ can be rewritten as:

$$-\delta\mu + \sum_{l=0}^{N} \binom{N}{l} p^l (1-p)^{N-l} [U_l' - \mu] \leq 0, \tag{43}$$

which is identical to the right hand side of (42) *except* for the weights $\left(\frac{l}{N}\right)$.

Note further that 1) $[f_l]$ implies that $[U_l' - \mu] \geq 0$ for all $l \neq 0$, and 2) $\left(\frac{l}{N}\right) < 1$ for $l < N$; $\left(\frac{l}{N}\right) = 1$ for $l = N$. This implies that the right hand side of (42) will be strictly less than the left hand side of (43) unless $U_l' - \mu = 0$ for $l < N$.

In other words, for catastrophe bond issuance will be strictly suboptimal unless consumers are fully insured in all states of the world *except* the one where *everyone* experiences a loss. In this scenario, the social planner is indifferent between using reinsurance policies alone and reinsurance policies in conjunction with catastrophe bonds.

C Proofs for Section 4.2

C.1 Derivation of (19)

The generalized asset allocation model produces the following first order conditions that analogize to the simpler model:

$$[B_i] \ : \ \sum_{\mathbf{s}\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_i^{s1}} \mu_\Theta(\theta)\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\left(\delta_B + \sum_{\mathbf{s}\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_i^{s1}} \mu_\Theta(\theta)\pi^\theta\right) \leq 0, \forall i \tag{44}$$

$$[c_i] \ : \ -\varphi_i + \mu = 0, \forall i \tag{45}$$

$$[I_i] \ : \ \sum_{\mathbf{s}\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_i^{s1}} \mu_\Theta(\theta) f_{\boldsymbol{\theta}}\left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta\right) - \sum_{\mathbf{s}\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_i^{s1}} f_{\boldsymbol{\theta}}\lambda_{\boldsymbol{\theta}} = 0, \forall i \tag{46}$$

$$[A] \ : \ -\delta_A\mu + \sum_{\mathbf{s}\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon^s} \lambda_{\boldsymbol{\theta}} = 0 \tag{47}$$

$$[f_{\boldsymbol{\theta}}] \ : \ \sum_{i\in\Gamma(\theta)} \mu_\Theta(\theta) I_i\left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta\right) - \lambda_\theta \sum_{i\in\Gamma(\boldsymbol{\theta})} I_i - \phi_{\boldsymbol{\theta}} = 0, \forall\theta \tag{48}$$

It is straightforward to derive the return on catastrophe bonds in this case, by following the outline of our earlier argument. First, we show that $\phi_\theta = 0$ for all θ. To see this, multiply $[I_i]$ by I_i and sum over i to obtain:

$$\sum_i\sum_{\mathbf{s}\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_i^{s1}} \mu_\Theta(\theta) f_{\boldsymbol{\theta}}\left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta\right) I_i - \sum_i\sum_{\mathbf{s}\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_i^{s1}} f_{\boldsymbol{\theta}}\lambda_{\boldsymbol{\theta}} I_i = 0$$

Next, multiply $[f_\theta]$ by f_θ and sum over θ to obtain:

$$\sum_{\theta\in\Theta}\sum_{i\in\Gamma(\theta)} \mu_\Theta(\theta) f_\theta I_i\left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta\right) - \sum_{\theta\in\Theta}\lambda_\theta f_\theta \sum_{i\in\Gamma(\boldsymbol{\theta})} I_i - \sum_{\theta\in\Theta} f_\theta\phi_{\boldsymbol{\theta}} = 0$$

The definitions of Γ and Υ imply that the double summation $\sum_{\theta\in\Theta}\sum_{i\in\Gamma(\theta)}$ is equivalent to $\sum_i\sum_{\mathbf{s}\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_i^{s1}}$. Therefore, it is the case that $\sum_{\theta\in\Theta} f_\theta\phi_{\boldsymbol{\theta}} = 0$, and that $\phi_{\boldsymbol{\theta}} = 0$ for all θ.

Using this result, and multiplying $[f_\theta]$ by f_θ, we obtain:

$$\sum_{i\in\Gamma(\theta)} \mu_\Theta(\theta) w_i^\theta f_\theta\left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta\right) - \lambda_\theta f_\theta = 0,$$

where

$$w_i^\theta = \frac{I_i}{\sum_{i\in\Gamma(\boldsymbol{\theta})} I_i}$$

Summing over $\theta \in \Theta^i$, we obtain:

$$\sum_{\theta\in\Theta^i} \mu_\Theta(\theta) f_\theta\left(\sum_{i\in\Gamma(\theta)} w_i^\theta\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta\right) = \sum_{\theta\in\Theta^i}\lambda_\theta f_\theta, \tag{49}$$

Dividing both sides of this equation by f_θ yields:

$$\sum_{\theta\in\Theta^i} \mu_\Theta(\theta) \left(\sum_{i\in\Gamma(\theta)} w_i^\theta \frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta \right) = \sum_{\theta\in\Theta^i} \lambda_\theta, \tag{50}$$

Subtracting equation 50 from 49 leads to:

$$\sum_{\theta\in\Theta^i} \mu_\Theta(\theta)(f_\theta - 1) \left(\sum_{i\in\Gamma(\theta)} w_i^\theta \frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta \right) = \sum_{\theta\in\Theta^i} \lambda_\theta(f_\theta - 1) \tag{51}$$

We now rewrite the marginal condition $[B_i]$ in terms of the Θ^i set notation:

$$R_i = \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\theta) \left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta \right) - \mu\delta_B \leq 0 \tag{52}$$

Subtracting the right-hand side of $[I_i]$ from the above equation results in:

$$R_i = \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\theta)(1 - f_\theta) \left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta \right) + \sum_{\theta\in\Theta^i} f_{\boldsymbol{\theta}}\lambda_{\boldsymbol{\theta}} - \mu\delta_B \leq 0 \tag{53}$$

or

$$R_i = \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\theta)(1 - f_\theta) \left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta \right) + \sum_{\theta\in\Theta^i} (f_{\boldsymbol{\theta}} - 1)\lambda_{\boldsymbol{\theta}} + \sum_{\theta\in\Theta^i} \lambda_\theta - \mu\delta_B \leq 0 \tag{54}$$

Substituting in from $[A]$ yields:

$$R_i = \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\theta)(1 - f_\theta) \left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \mu\pi^\theta \right) + \sum_{\theta\in\Theta^i} (f_{\boldsymbol{\theta}} - 1)\lambda_{\boldsymbol{\theta}} - \sum_{\theta\notin\Theta^i} \lambda_\theta + \mu(\delta_A - \delta_B) \leq 0 \tag{55}$$

Substituting in equation 51 transforms this into:

$$R_i = \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\theta)(1 - f_\theta) \left(\frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} - \sum_{i\in\Gamma(\theta)} w_i^\theta \frac{\partial U_i^{\boldsymbol{\theta}}}{\partial W} \right) - \sum_{\theta\notin\Theta^i} \lambda_\theta + \mu(\delta_A - \delta_B) \leq 0 \tag{56}$$

This is the final expression for the bond return.

C.2 Characterization of Solution without Frictional Costs

This section presents four lemmas in support of the claims in Section 4.2. Specifically, the intention is to show that, in the absence of frictional costs, 1) full insurance is optimal, and 2) given full insurance, the division of risk transfer between reinsurance polices and catastrophe bonds is irrelevant. The approach is as follows.

The first lemma facilitates later work. It shows that, for each consumer, the sum of promised indemnification under insurance policies and catastrophe bond principal in any solution will be at least equal to the loss, allowing us to rule out solutions where some consumers have less than full coverage.

The second lemma shows that, if the optimal solution features exactly full coverage for each consumer (in the sense of the sum of promised indemnification under insurance policies and catastrophe bond principal being exactly equal to the loss), there will be no default: In other words, full coverage means full insurance.

The third lemma shows the indeterminacy of risk transfer: If consumers have exactly full coverage (and hence, by the second lemma, full insurance), it does not matter how that coverage is delivered—a continuum of solutions with exactly full coverage (but different splits of coverage between cat bonds and insurance) exist.

The fourth lemma shows that, in any case where a solution exists where a consumer has greater than full coverage (i.e., the sum of promised indemnity and catastrophe bond principal is *greater* than the loss), an equivalent solution can be constructed with exactly full coverage. This completes the proof of the "full insurance is always optimal." By the first lemma, we know that only solutions with exactly full coverage and greater than full coverage need be considered. By the fourth lemma, we know that any solution with greater than fully coverage can be replicated by a solution with exactly full coverage. By the second lemma, we know that a solution with full coverage also features no default (i.e., full insurance).

Likewise, the fourth lemma completes the proof of the second claim—that the division of risk transfer between insurance policies and catastrophe bonds is indeterminate in the absence of frictional costs. By the reasoning in the preceding paragraph, we know that solutions with exactly full coverage and full insurance are always optimal in the absence of frictional costs. By the third lemma, we know that the division of risk transfer between insurance policies and catastrophe bonds is indeterminate in solutions with exactly full coverage and full insurance, in the absence of frictional costs.

Lemma. *Consider the maximization of (14) subject to the constraints below it, and assume at least one solution exists. Then there exists a solution* $\hat{A}, \{\hat{B}_i\}, \{\hat{c}_i\}, \{\hat{I}_i\}, \{\hat{f}_{\boldsymbol{\theta}}\}, \{\hat{W}_{is}\}$ *with* $\hat{I}_i + \hat{B}_i \geq L_i$*, for all* i.

Proof:

Suppose not. Then there is a solution with $\hat{I}_k + \hat{B}_k < L_k$ for some $k \in \{1, ..., N\}$. Consider an alternative set of choices:

$\tilde{A} = \hat{A}$.

$\tilde{I}_i = \hat{I}_i$, for all i.

$\tilde{f}_{\boldsymbol{\theta}} = \hat{f}_{\boldsymbol{\theta}}$, for all $\boldsymbol{\theta}$.

$\tilde{B}_i = \hat{B}_i$, for all $i \neq k$; $\tilde{B}_k = \hat{B}_k + \left(L_k - (\hat{I}_k + \hat{B}_k)\right)$.

$\tilde{c}_i = \hat{c}_i$, for all $i \neq k$; $\tilde{c}_k = \hat{c}_k + \sum_{\boldsymbol{\theta} \in \Theta^k} \mu_\Theta(\boldsymbol{\theta}) \pi^\theta \left(\tilde{B}_k - \hat{B}_k\right)$

$\tilde{W}_{is} = \hat{W}_{is}$, $\forall i \neq k$, $\forall s$; $\tilde{W}_{ks} = \hat{W}_{ks} - \sum_{\boldsymbol{\theta} \in \{\boldsymbol{\theta}: \boldsymbol{\theta}(N+1)=s, \boldsymbol{\theta} \in \Theta^k\}} \mu_\Theta(\boldsymbol{\theta}) \left(\tilde{B}_k - \hat{B}_k\right)$, $\forall s$

We first verify that these new choices satisfy all constraints. First, note that

$$
\begin{aligned}
& \sum_{\boldsymbol{\theta}\in\Theta}\left(\mu_\Theta(\boldsymbol{\theta})\pi^\theta \tilde{f}_{\boldsymbol{\theta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} \tilde{I}_i\right) + \sum_i \left(\sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta\right)\tilde{B}_i \quad = \\
& \sum_{\boldsymbol{\theta}\in\Theta}\left(\mu_\Theta(\boldsymbol{\theta})\pi^\theta \hat{f}_{\boldsymbol{\theta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} \hat{I}_i\right) + \sum_i \left(\sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta\right)\hat{B}_i + \sum_{\boldsymbol{\theta}\in\Theta^k}\mu_\Theta(\boldsymbol{\theta})\pi^\theta\left(\tilde{B}_k - \hat{B}_k\right) \quad \leq \\
& \sum \hat{c}_i + \sum_{\boldsymbol{\theta}\in\Theta^k}\mu_\Theta(\boldsymbol{\theta})\pi^\theta\left(\tilde{B}_k - \hat{B}_k\right) \quad = \quad \sum \tilde{c}_i.
\end{aligned}
$$

So $[\mu]$ is satisfied.

Next, since $\tilde{A} = \hat{A}$, $\tilde{I}_i = \hat{I}_i$ (for all i), and $\tilde{f}_{\boldsymbol{\theta}} = \hat{f}_{\boldsymbol{\theta}}$ (for all $\boldsymbol{\theta}$), the new choices satisfy $[\lambda_{\boldsymbol{\theta}}]$ and $[\phi_{\boldsymbol{\theta}}]$.

Next, note that $[\varphi_i]$ is obviously satisfied for $i \neq k$. For $i = k$, note that:

$$
\begin{aligned}
\sum_{s\in\Psi}\pi_s \tilde{W}_{ks} &= \sum_{s\in\Psi}\pi_s\left(\hat{W}_{ks} - \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}}\mu_\Theta(\boldsymbol{\theta})\left(\tilde{B}_k - \hat{B}_k\right)\right) \\
&\leq W - \hat{c}_k - \sum_{s\in\Psi}\pi_s \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}}\mu_\Theta(\boldsymbol{\theta})\left(\tilde{B}_k - \hat{B}_k\right) = W - \tilde{c}_k,
\end{aligned}
$$

so $[\varphi_k]$ is satisfied.

The non-negativity constraints are obviously satisfied in all cases.

We next verify that the proposed alternative set of choices improves the objective function. Since $\tilde{I}_i = \hat{I}_i$ (for all i), $\tilde{B}_i = \hat{B}_i$ (for all $i \neq k$), $\tilde{c}_i = \hat{c}_i$ (for all $i \neq k$), $\tilde{W}_{is} = \hat{W}_{is}$ (for all $i \neq k$), and $\tilde{f}_{\boldsymbol{\theta}} = \hat{f}_{\boldsymbol{\theta}}$ (for all $\boldsymbol{\theta}$), it is obvious that the alternative choices yield identical levels of expected utility for all consumers other than consumer k.

Recall the definitions $\Upsilon_k^{s0} \equiv \{\boldsymbol{\theta} : \boldsymbol{\theta} \notin \Theta^k, \boldsymbol{\theta}(N+1) = s\}$ and $\Upsilon_k^{s1} \equiv \{\boldsymbol{\theta} : \boldsymbol{\theta} \in \Theta^k, \boldsymbol{\theta}(N+1) = s\}$, noting further that

$$
\Upsilon^s \equiv \Upsilon_k^{s0} \cup \Upsilon_k^{s1} = \{\boldsymbol{\theta} : \boldsymbol{\theta}(N+1) = s\}.
$$

For consumer k,

$$
\begin{aligned}
E\tilde{U}_k &= \sum_{s\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}}\mu_\Theta(\boldsymbol{\theta})U_k\left(\tilde{W}_{ks} - L_k + \tilde{f}_{\boldsymbol{\theta}}\tilde{I}_k + \tilde{B}_k\right) + \sum_{s\in\Psi}\sum_{\boldsymbol{\theta}\notin\Upsilon_k^{s0}}\mu_\Theta(\boldsymbol{\theta})U_k\left(\tilde{W}_{ks}\right) \\
E\hat{U}_k &= \sum_{s\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}}\mu_\Theta(\boldsymbol{\theta})U_k\left(\hat{W}_{ks} - L_k + \hat{f}_{\boldsymbol{\theta}}\hat{I}_k + \hat{B}_k\right) + \sum_{s\in\Psi}\sum_{\boldsymbol{\theta}\notin\Upsilon_k^{s0}}\mu_\Theta(\boldsymbol{\theta})U_k\left(\hat{W}_{ks}\right)
\end{aligned}
$$

The difference $E\tilde{U}_k - E\hat{U}_k$ reflects the increase in utility when moving from the "solution" to the utility level obtained with the alternative set of choices. It can be written as:

$$\sum_{s\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k\left(\tilde{W}_{ks} - L_k + \tilde{f}_{\boldsymbol{\theta}}\tilde{I}_k + \tilde{B}_k\right) - U_k\left(\hat{W}_{ks} - L_k + \hat{f}_{\boldsymbol{\theta}}\hat{I}_k + \hat{B}_k\right)\right] +$$
$$\sum_{s\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k\left(\tilde{W}_{ks}\right) - U_k\left(\hat{W}_{ks}\right)\right]$$

If this difference is weakly positive, then the alternative set of choices constitutes an additional solution, and we have our desired contradiction.

If security markets are complete with respect to consumer k's loss exposure (i.e., if, for every s, either $\Upsilon_k^{s0} = \emptyset$ or $\Upsilon_k^{s1} = \emptyset$), then it is trivial to show that this difference is zero. If security markets are incomplete with respect to consumer k's loss exposure, then there will be some nonempty subset $Z \subset \Psi$ such that:

$$Z = \{s : \Upsilon_k^{s0} \neq \emptyset, \Upsilon_k^{s1} \neq \emptyset\}.$$

For every $s \in Z$,

$$\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k\left(\tilde{W}_{ks} - L_k + \tilde{f}_{\boldsymbol{\theta}}\tilde{I}_k + \tilde{B}_k\right) - U_k\left(\hat{W}_{ks} - L_k + \hat{f}_{\boldsymbol{\theta}}\hat{I}_k + \hat{B}_k\right)\right] +$$
$$\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k\left(\tilde{W}_{ks}\right) - U_k\left(\hat{W}_{ks}\right)\right]$$

may be rewritten as:

$$\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_k\left(\tilde{W}_{ks} - L_k + \tilde{f}_{\boldsymbol{\theta}}\tilde{I}_k + \tilde{B}_k\right) + \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_k\left(\tilde{W}_{ks}\right) -$$
$$\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_k\left(\hat{W}_{ks} - L_k + \hat{f}_{\boldsymbol{\theta}}\hat{I}_k + \hat{B}_k\right) + \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_k\left(\hat{W}_{ks}\right)$$

and further as:

$$\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_k\left(\hat{W}_{ks} - \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta})\left(\tilde{B}_k - \hat{B}_k\right) - L_k + \tilde{f}_{\boldsymbol{\theta}}\tilde{I}_k + \left(\tilde{B}_k - \hat{B}_k\right) + \tilde{B}_k\right)$$
$$+ \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_k\left(\hat{W}_{ks} - \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta})\left(\tilde{B}_k - \hat{B}_k\right)\right) -$$
$$\left(\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_k\left(\hat{W}_{ks} - L_k + \hat{f}_{\boldsymbol{\theta}}\hat{I}_k + \hat{B}_k\right) + \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_k\left(\hat{W}_{ks}\right)\right),$$

which is greater than zero, since U_k is concave and

$$\hat{f}_{\boldsymbol{\theta}}\hat{I}_k + \hat{B}_k < \tilde{f}_{\boldsymbol{\theta}}\tilde{I}_k + \left(\tilde{B}_k - \hat{B}_k\right) + \tilde{B}_k \leq L_k,\ \forall\boldsymbol{\theta}.$$

QED.

Lemma. *Let* $\hat{A}, \{\hat{B}_i\}, \{\hat{c}_i\}, \{\hat{I}_i\}, \{\hat{f}_{\boldsymbol{\theta}}\}, \{\hat{W}_{is}\}$ *be choices that solve (14) subject to the constraints below it, and suppose* $\hat{I}_i + \hat{B}_i = L$*, for all* i*. Then there exists a solution* $\tilde{A}, \{\tilde{B}_i\}, \{\tilde{c}_i\}, \{\tilde{I}_i\}, \{\tilde{f}_{\boldsymbol{\theta}}\}, \{\tilde{W}_{is}\}$ *with 1)* $\tilde{I}_i + \tilde{B}_i = L_i$*, for all* i*, and 2)* $\tilde{f}_{\boldsymbol{\theta}} = 1$*, for all* $\boldsymbol{\theta}$ *where* $\Gamma(\boldsymbol{\theta})$ *is nonempty and* $\hat{I}_i > 0$ *for some* $i \in \Gamma(\boldsymbol{\theta})$*.*

Proof:

Suppose not. Then the set $\Delta = \{\boldsymbol{\theta} : \hat{f}_{\boldsymbol{\theta}} < 1\}$ is nonempty and $\hat{I}_i > 0$ for some $i \in \Gamma(\boldsymbol{\theta})$ for some $\boldsymbol{\theta} \in \boldsymbol{\Delta}$.

Let $\Phi \equiv \{i : \hat{I}_i > 0,\ i \in \Gamma(\boldsymbol{\theta})$ for some $\boldsymbol{\theta} \in \boldsymbol{\Delta}\}$. Consider an alternative set of choices as follows:

$\tilde{f}_{\boldsymbol{\theta}} = 1,\ \forall\boldsymbol{\theta}$.

$\tilde{I}_i = \hat{I}_i$, for all i.

$\tilde{B}_i = \hat{B}_i$, for all i

$\tilde{A} = \max\left[\hat{A},\ \max_{\boldsymbol{\theta}\in\boldsymbol{\Delta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} \tilde{I}_i\right]$

$\tilde{c}_i = \hat{c}_i + \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left(\tilde{f}_{\boldsymbol{\theta}} - \hat{f}_{\boldsymbol{\theta}}\right) \tilde{I}_i,\ \forall i.$

$\tilde{W}_{is} = \hat{W}_{is} - \sum_{\boldsymbol{\theta}\in\{\boldsymbol{\theta}:\boldsymbol{\theta}\in\boldsymbol{\Theta}^{\mathbf{i}},\boldsymbol{\theta}(N+1)=s\}} \mu_\Theta(\boldsymbol{\theta}) \left(\tilde{f}_{\boldsymbol{\theta}} - \hat{f}_{\boldsymbol{\theta}}\right) \tilde{I}_i,\ \forall s,\ \forall i.$

We first verify that the constraints are satisfied.

Starting with $[\mu]$, note that

$$\begin{aligned}
\sum \tilde{c}_i &= \sum_i \left(\hat{c}_i + \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left(\tilde{f}_{\boldsymbol{\theta}} - \hat{f}_{\boldsymbol{\theta}}\right) \tilde{I}_i\right) \\
&= \sum_{\boldsymbol{\theta}\in\Theta} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \hat{f}_{\boldsymbol{\theta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} \hat{I}_i + \sum_i \left(\sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta\right) \hat{B}_i + \sum_i \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left(\tilde{f}_{\boldsymbol{\theta}} - \hat{f}_{\boldsymbol{\theta}}\right) \tilde{I}_i \\
&= \sum_{\boldsymbol{\theta}\in\Theta} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \tilde{f}_{\boldsymbol{\theta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} \tilde{I}_i + \sum_i \left(\sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta\right) \tilde{B}_i.
\end{aligned}$$

By construction of $\tilde{A}$, it is clear that $[\lambda_{\boldsymbol{\theta}}]$ is satisfied. Likewise, since $\tilde{f}_{\boldsymbol{\theta}} = 1,\ \forall\boldsymbol{\theta}$, $[\phi_{\boldsymbol{\theta}}]$ is satisfied.

For the individual wealth constraint $[\varphi_i]$,

$$\begin{aligned}
\sum_s \pi_s \tilde{W}_{is} &= \sum_s \pi_s \hat{W}_{is} - \sum_s \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left(\tilde{f}_{\boldsymbol{\theta}} - \hat{f}_{\boldsymbol{\theta}}\right) \tilde{I}_i \\
&\leq W - \hat{c}_i - \sum_s \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left(\tilde{f}_{\boldsymbol{\theta}} - \hat{f}_{\boldsymbol{\theta}}\right) \tilde{I}_i \\
&= W - \tilde{c}_i.
\end{aligned}$$

The non-negativity constraints are obviously satisfied. We now must verify that the alternative choices improve the objective function.

Evidently, utility is unchanged for $i \notin \Phi$, so we are left to verify that utility improves for $i \in \Phi$. We thus consider some consumer k (with $k \in \Phi$). To verify that utility improves with the new choices, we focus on the utility levels under the alternative and original choices, as in:

$$\begin{aligned}
E\tilde{U}_k &= \sum_s \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_k \left(\tilde{W}_{ks} - L_k + \tilde{f}_{\boldsymbol{\theta}} \tilde{I}_k + \tilde{B}_k \right) + \sum_s \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_k \left(\tilde{W}_{ks} \right) \\
E\hat{U}_k &= \sum_s \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_k \left(\hat{W}_k^{\boldsymbol{\theta}} - L_k + \hat{f}_{\boldsymbol{\theta}} \hat{I}_k + \hat{B}_k \right) + \sum_s \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_k \left(\hat{W}_k^{\boldsymbol{\theta}} \right)
\end{aligned}$$

and examine the difference $E\tilde{U}_k - E\hat{U}_k$.

If security markets are complete with respect to consumer k's loss exposure (i.e., if, for every s, either $\Upsilon_k^{s0} = \emptyset$ or $\Upsilon_k^{s1} = \emptyset$), then it can be shown that this difference is zero—since security markets can replicate the payoffs from insurance policies for consumer k. Similarly, when security markets are incomplete, but the sets $Z \equiv \{s : \Upsilon_k^{s0} \neq \emptyset, \Upsilon_k^{s1} \neq \emptyset\}$ and $X = \{\boldsymbol{\theta}(N+1) : \boldsymbol{\theta} \in \Delta\}$ do not intersect, the difference will be zero for the same reason.

On the other hand, if $Z \cap X \neq \emptyset$, then consumer k cannot replicate insurance policy payoffs with securities. For every s,

$$\begin{aligned}
&\sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k \left(\tilde{W}_{ks} - L_k + \tilde{f}_{\boldsymbol{\theta}} \tilde{I}_k + \tilde{B}_k \right) - U_k \left(\hat{W}_{ks} - L_k + \hat{f}_{\boldsymbol{\theta}} \hat{I}_k + \hat{B}_k \right) \right] + \\
&\sum_{\boldsymbol{\theta} \in \Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k \left(\tilde{W}_k^{\boldsymbol{\theta}} \right) - U_k \left(\hat{W}_k^{\boldsymbol{\theta}} \right) \right]
\end{aligned}$$

may be rewritten as:

$$\begin{aligned}
&\sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_k \left(\hat{W}_{ks} - \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left(1 - \hat{f}_{\boldsymbol{\theta}} \right) \hat{I}_k \right) + \\
&\sum_{\boldsymbol{\theta} \in \Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_k \left(\hat{W}_{ks} - \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left(1 - \hat{f}_{\boldsymbol{\theta}} \right) \hat{I}_k \right) - \\
&\sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) U_k \left(\hat{W}_{ks} - \left(1 - \hat{f}_{\boldsymbol{\theta}} \right) \hat{I}_k \right) + \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) U_k \left(\hat{W}_{ks} \right)
\end{aligned}$$

which is greater than or equal to zero, since U_k is concave, with strict inequality when $s \in Z \cap X$.

QED.

Lemma. *Let* $\hat{A}, \{\hat{B}_i\}, \{\hat{c}_i\}, \{\hat{I}_i\}, \{\hat{f}_{\boldsymbol{\theta}}\}, \{\hat{W}_{is}\}$ *be choices that solve (14) subject to the constraints below it, and* $\hat{I}_i + \hat{B}_i = L$*, for all* i*. Then, an equivalent solution yielding 1) the same value for the overall objective function and 2) the same values for the utility levels of each of the individual consumers, can be constructed as* $\hat{A}, \{\tilde{B}_i\}, \{\hat{c}_i\}, \{\tilde{I}_i\}, \{\hat{f}_{\boldsymbol{\theta}}\}, \{\hat{W}_{is}\}$*, where* $\{\tilde{B}_i\}$ *and* $\{\tilde{I}_i\}$ *are any alternative sets of feasible choices satisfying the nonnegativity constraints and* $\tilde{I}_i + \tilde{B}_i = L_i$*, for all* i*, and* $\hat{A} = \max_{\boldsymbol{\theta}} \sum_{\Gamma(\boldsymbol{\theta})} \tilde{I}_i$.

Proof:

The preceding lemma shows that for any solution with $\hat{I}_i + \hat{B}_i = L_i$, for all i, an equivalent solution can be constructed as:

$\check{f}_{\boldsymbol{\theta}} = 1, \ \forall \boldsymbol{\theta}$.

$\check{I}_i = \hat{I}_i$, for all i.

$\check{B}_i = \hat{B}_i$, for all i

$\check{A} = \max \left[\hat{A}, \ \max_{\boldsymbol{\theta} \in \boldsymbol{\Delta}} \sum_{i \in \Gamma(\boldsymbol{\theta})} \check{I}_i\right]$

$\check{c}_i = \hat{c}_i + \sum_{\boldsymbol{\theta} \in \Theta^i} \mu_\Theta(\boldsymbol{\theta}) \pi^\theta \left(\check{f}_{\boldsymbol{\theta}} - \hat{f}_{\boldsymbol{\theta}}\right) \check{I}_i, \ \forall i.$

$\check{W}_{is} = \hat{W}_{is} - \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left(\check{f}_{\boldsymbol{\theta}} - \hat{f}_{\boldsymbol{\theta}}\right) \check{I}_i, \ \forall s, \ \forall i.$

Then, given any sets $\{\tilde{B}_i\}$ and $\{\tilde{I}_i\}$ that satisfy $\tilde{I}_i + \tilde{B}_i = L_i$, $\tilde{I}_i \geq 0$, and $\tilde{B}_i \geq 0$, for all i, we construct the remaining choice variables as:

$\tilde{f}_{\boldsymbol{\theta}} = 1, \ \forall \boldsymbol{\theta}$.

$\tilde{A} = \max_{\boldsymbol{\theta}} \sum_{i \in \Gamma(\boldsymbol{\theta})} \tilde{I}_i$

$\tilde{c}_i = \check{c}_i, \ \forall i.$

$\tilde{W}_{is} = \check{W}_{is}, \ \forall s, \ \forall i.$

We now must show that the new choices satisfy all constraints and (weakly) improve the objective function. Starting with constraint $[\mu]$, in the absence of frictional costs, we have:

$$\sum \tilde{c}_i = \sum \check{c}_i \geq \sum_{\boldsymbol{\theta} \in \Theta} \left(\mu_\Theta(\boldsymbol{\theta}) \pi^\theta f_{\boldsymbol{\theta}} \sum_{i \in \Gamma(\boldsymbol{\theta})} \check{I}_i \right) + \sum_i \left(\sum_{\boldsymbol{\theta} \in \Theta^i} \mu_\Theta(\boldsymbol{\theta}) \pi^\theta \right) \check{B}_i$$

Since $f_{\boldsymbol{\theta}} = 1$, we may rewrite as:

$$\sum \tilde{c}_i = \sum \check{c}_i \geq \sum_i \left(\sum_{\boldsymbol{\theta} \in \Theta^i} \mu_\Theta(\boldsymbol{\theta}) \pi^\theta \right) \left(\check{B}_i + \check{I}_i\right)$$

But, since

$$\check{B}_i + \check{I}_i = \tilde{I}_i + \tilde{B}_i = L,$$

$$\sum \tilde{c}_i \geq \sum_i \left(\sum_{\boldsymbol{\theta} \in \Theta^i} \mu_\Theta(\boldsymbol{\theta}) \pi^\theta \right) \left(\tilde{B}_i + \tilde{I}_i\right),$$

so the constraint is satisfied.

For constraint $[\lambda_{\boldsymbol{\theta}}]$, note that:

$$\tilde{A} = \max_{\boldsymbol{\theta}} \sum_{i \in \Gamma(\boldsymbol{\theta})} \tilde{I}_i,$$

which is obviously greater than or equal to $f_{\boldsymbol{\theta}} \sum_{i \in \Gamma(\boldsymbol{\theta})} \tilde{I}_i$ for all $\boldsymbol{\theta}$, so the constraint is satisfied.

Constraints $[\phi_{\boldsymbol{\theta}}]$ and $[\varphi_i]$ are obviously satisfied, since $\tilde{f}_{\boldsymbol{\theta}} = 1$, $\tilde{c}_i = \check{c}_i$, and $\tilde{W}_{is} = \check{W}_{is}$.

Since 1) state contingent wealth allocations are unchanged under the new choices ($\tilde{W}_{is} = \check{W}_{is}$, for all s and for all i), and 2) the total of insurance policy and cat bond recoveries are unchanged for every consumer in every state of the world where that consumer experiences a loss:

$$f_{\boldsymbol{\theta}} \tilde{I}_i + \tilde{B}_i = \tilde{I}_i + \tilde{B}_i = \check{I}_i + \check{B}_i = f_{\boldsymbol{\theta}} \check{I}_i + \check{B}_i, \ \forall i, \boldsymbol{\theta} \in \Theta^i,$$

it is evident that utility is unchanged for every consumer. Thus, the objective function is unchanged and the alternative choices proposed must constitute an additional solution.

QED.

Lemma. *Let $\hat{A}, \{\hat{B}_i\}, \{\hat{c}_i\}, \{\hat{I}_i\}, \{\hat{f}_{\boldsymbol{\theta}}\}, \{\hat{W}_{is}\}$ be choices that solve (14) subject to the constraints below it, and suppose $\hat{I}_i + \hat{B}_i > L$, for at least one i. Then there exists a "full coverage, full insurance" solution yielding the same value for the overall objective function that can be constructed with $\{\tilde{B}_i\}$ and $\{\tilde{I}_i\}$ as alternative sets of feasible choices satisfying the nonnegativity constraints and $\tilde{I}_i + \tilde{B}_i = L_i$, for all i, $\tilde{f}_{\boldsymbol{\theta}} = 1 \ \forall \boldsymbol{\theta}$, and $\tilde{A} = \max_{\boldsymbol{\theta}} \sum_{i \in \Gamma(\boldsymbol{\theta})} \tilde{I}_i$.*

Proof:

From the preceding lemma, we know that all solutions with $\tilde{I}_i + \tilde{B}_i = L_i$, for all i, and $\tilde{f}_{\boldsymbol{\theta}} = 1$, $\forall \boldsymbol{\theta}$, are equivalent. Given sets $\{\tilde{B}_i\}$ and $\{\tilde{I}_i\}$ satisfying these conditions, construct the remaining choice variables as:

$\tilde{A} = \max_{\boldsymbol{\theta}} \sum_{i \in \Gamma(\boldsymbol{\theta})} \tilde{I}_i$

$\tilde{f}_{\boldsymbol{\theta}} = 1, \ \forall \boldsymbol{\theta}$

$\tilde{c}_i = \hat{c}_i + \sum_{\boldsymbol{\theta} \in \Theta^i} \mu_\Theta(\boldsymbol{\theta}) \pi^\theta \left[\left(\tilde{f}_{\boldsymbol{\theta}} \tilde{I}_i - \hat{f}_{\boldsymbol{\theta}} \hat{I}_i \right) + \left(\tilde{B}_i - \hat{B}_i \right) \right]$

$\tilde{W}_{is} = \hat{W}_{is} - \sum_{\boldsymbol{\theta} \in \Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left[\left(\tilde{f}_{\boldsymbol{\theta}} \tilde{I}_i - \hat{f}_{\boldsymbol{\theta}} \hat{I}_i \right) + \left(\tilde{B}_i - \hat{B}_i \right) \right], \ \forall s, \ \forall i$

We start by verifying that these alternative choices satisfy the constraints. First, rewrite the original constraint as:

$$\sum \hat{c}_i \geq \sum_i \sum_{\boldsymbol{\theta} \in \Theta^i} \mu_\Theta(\boldsymbol{\theta}) \pi^\theta \left(\hat{f}_{\boldsymbol{\theta}} \hat{I}_i + \hat{B}_i \right).$$

Furthermore,

$$
\begin{aligned}
\sum \tilde{c}_i &= \sum \hat{c}_i + \sum_i \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left[\left(\tilde{f}_{\boldsymbol{\theta}}\tilde{I}_i - \hat{f}_{\boldsymbol{\theta}}\hat{I}_i\right) + \left(\tilde{B}_i - \hat{B}_i\right)\right] \\
&\geq \sum_i \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left(\hat{f}_{\boldsymbol{\theta}}\hat{I}_i + \hat{B}_i\right) + \sum_i \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left[\left(\tilde{f}_{\boldsymbol{\theta}}\tilde{I}_i - \hat{f}_{\boldsymbol{\theta}}\hat{I}_i\right) + \left(\tilde{B}_i - \hat{B}_i\right)\right] \\
&\geq \sum_i \sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left(\tilde{f}_{\boldsymbol{\theta}}\tilde{I}_i + \tilde{B}_i\right) \\
&= \sum_{\boldsymbol{\theta}\in\Theta} \left(\mu_\Theta(\boldsymbol{\theta})\pi^\theta \tilde{f}_{\boldsymbol{\theta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} \tilde{I}_i\right) + \sum_i \left(\sum_{\boldsymbol{\theta}\in\Theta^i} \mu_\Theta(\boldsymbol{\theta})\pi^\theta\right) \tilde{B}_i
\end{aligned}
$$

so $[\mu]$ is satisfied.

Next, since $\tilde{A} = \max_{\boldsymbol{\theta}} \sum_{i\in\Gamma(\boldsymbol{\theta})} \tilde{I}_i$ and $\tilde{f}_{\boldsymbol{\theta}} = 1$ for all $\boldsymbol{\theta}$, $[\lambda_{\boldsymbol{\theta}}]$ and $[\phi_{\boldsymbol{\theta}}]$ are obviously satisfied.

For $[\varphi_i]$, note that, for each i,

$$
\begin{aligned}
\sum_{s\in\Psi} \pi_s \tilde{W}_{is} &= \sum_{s\in\Psi} \pi_s \hat{W}_{is} - \sum_{s\in\Psi} \pi_s \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left[\left(\tilde{f}_{\boldsymbol{\theta}}\tilde{I}_i - \hat{f}_{\boldsymbol{\theta}}\hat{I}_i\right) + \left(\tilde{B}_i - \hat{B}_i\right)\right] \\
&\leq W - \hat{c}_i - \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta})\pi^\theta \left[\left(\tilde{f}_{\boldsymbol{\theta}}\tilde{I}_i - \hat{f}_{\boldsymbol{\theta}}\hat{I}_i\right) + \left(\tilde{B}_i - \hat{B}_i\right)\right] \\
&= W - \tilde{c}_i,
\end{aligned}
$$

so the constraints are evidently satisfied.

By definition, the non-negativity constraints are satisfied.

It remains to show that the objective function (weakly) improves under the alternative choices. For any consumer k, the difference between utility achieved under the alternative choices and the original choices can be expressed as:

$$
\begin{aligned}
&\sum_{s\in\Psi} \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k\left(\tilde{W}_{ks} - L_k + \tilde{f}_{\boldsymbol{\theta}}\tilde{I}_k + \tilde{B}_k\right) - U_k\left(\hat{W}_{ks} - L_k + \hat{f}_{\boldsymbol{\theta}}\hat{I}_k + \hat{B}_k\right)\right] + \\
&\sum_{s\in\Psi} \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k\left(\tilde{W}_{ks}\right) - U_k\left(\hat{W}_{ks}\right)\right].
\end{aligned}
$$

Simplifying (recall that $\tilde{f}_{\boldsymbol{\theta}}\tilde{I}_k + \tilde{B}_k$ always equals L) yields:

$$
\begin{aligned}
&\sum_{s\in\Psi} \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k\left(\tilde{W}_{ks}\right) - U_k\left(\hat{W}_{ks} - L_k + \hat{f}_{\boldsymbol{\theta}}\hat{I}_k + \hat{B}_k\right)\right] + \\
&\sum_{s\in\Psi} \sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}} \mu_\Theta(\boldsymbol{\theta}) \left[U_k\left(\tilde{W}_{ks}\right) - U_k\left(\hat{W}_{ks}\right)\right],
\end{aligned}
$$

or:

$$\sum_{s\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}}\mu_\Theta(\boldsymbol{\theta})\left[U_k\left(\hat{W}_{ks}-\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}}\mu_\Theta(\boldsymbol{\theta})\left[L_k-(\hat{f}_{\boldsymbol{\theta}}\hat{I}_k+\hat{B}_k)\right]\right)-U_k\left(\hat{W}_{ks}-L_k+\hat{f}_{\boldsymbol{\theta}}\hat{I}_k+\hat{B}_k\right)\right]$$
$$+\sum_{s\in\Psi}\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s0}}\mu_\Theta(\boldsymbol{\theta})\left[U_k\left(\hat{W}_{ks}-\sum_{\boldsymbol{\theta}\in\Upsilon_k^{s1}}\mu_\Theta(\boldsymbol{\theta})\left[L_k-(\hat{f}_{\boldsymbol{\theta}}\hat{I}_k+\hat{B}_k)\right]\right)-U_k\left(\hat{W}_{ks}\right)\right]$$

Since U_k is concave, it follows that this difference is weakly greater than zero.

QED.

References

[1] Borch, K. (1962), "Equilibrium in a Reinsurance Market," *Econometrica* 30, 424-44.

[2] Bouriaux, S., and Scott, W.L. (2004), "Capital Market Solutions to Terrorism Risk Coverage: A Feasibility Study," *Journal of Risk Finance* 5(4).

[3] Brandts, S., and Laux, C. (2005), "ART versus Reinsurance: The Disciplining Effect of Information Sensitivity," *Paper presented at the 2006 NBER Insurance Workshop.*

[4] Cox, L.A., and Rudd, E.A. (1991) "Book Versus Market Underwriting Betas," *Journal of Risk and Insurance* 58, 312-21.

[5] Cummins, J.D., and Harrington, S.E. (1985), "Property-Liability Insurance Rate Regulation: Estimation of Underwriting Betas Using Quarterly Profit Data," *Journal of Risk and Insurance* 52, 16-43.

[6] Cummins, J.D., Phillips, R.D., and Allen, F. (1998), "Financial Pricing of Insurance in the Multiple-Line Insurance Company," *Journal of Risk and Insurance* 65, 597-636.

[7] Doherty, N.A., and Mahul, O. (2001), "Mickey Mouse and Moral Hazard: Uninformative but Correlated Triggers," working paper, Wharton.

[8] Doherty, N.A., and Schlesinger, H. (1990), "Rational Insurance Purchasing: Consideration of Contract Nonperformance," *Quarterly Journal of Economics* 105, 243-53.

[9] Duffie, D. (1992), *Dynamic Asset Pricing Theory*, Princeton: Princeton University Press.

[10] Froot, K.A. (ed.) (1999), *The Financing of Catastrophe Risk.* Chicago and London: University of Chicago Press.

[11] Froot, K.A. (2001), "The Market for Catastrophe Risk: A Clinical Examination," *Journal of Financial Economics* 60, 529-71.

[12] Hall, R.M. (2001), "Cedents' Claims to Reinsurance Recoverables and Priority of Distribution," *Mealey's Litigation Report: Reinsurance* 12, No. 13.

[13] Hall, R.M. (2000), "Drawing Down Letters of Credit in an Insurer Receivership Context," *Mealey's Litigation Report: Insurance Insolvency* 11, No. 21.

[14] Hoyt, R.E., and McCullough, K.A. (1999), "Catastrophe Insurance Options: Are They Zero-Beta Assets?" *Journal of Insurance Issues* 22, 147-63.

[15] Jaffee, D. (2006), "Monoline Restrictions, with Applications to Mortgage Insurance and Title Insurance," *Review of Industrial Organization,* forthcoming.

[16] Litzenberger, R.H., Beaglehole, D.R., and Reynolds, C.E. (1996), "Assessing Catastrophe Reinsurance-Linked Securities as a New Asset Class," *Journal of Portfolio Management* (Special Issue - December 1996), 76-86.

[17] Mahul, O., and Wright, B.D. (2003), "Efficient Risk Sharing within a Catastrophe Insurance Pool," *Paper presented at the 2003 NBER Insurance Workshop.*

[18] Mayers, D., and Smith, C.W. (1983), "The Interdependence of Individual Portfolio Decisions and the Demand for Insurance," *Journal of Political Economy* 91, 304-11.

[19] Murray, A. (2006), "Reinsurance Side-Cars: Going Along for the Ride," *Special Comment*, Moody's Investor Service.

[20] Myers, S.C. and Read, J.A. (2001), "Capital Allocation for Insurance Companies," *Journal of Risk and Insurance* 68, 545-80.

[21] Niehaus, G. (2002), "The Allocation of Catastrophe Risk," *Journal of Banking and Finance* 26, 585-96.

[22] Woo, G. (2004), "A Catastrophe Bond Niche: Multiple Event Risk," Paper presented at the 2004 meeting of the NBER Insurance Project Group.

[23] Zanjani, G. (2002), "Pricing and Capital Allocation in Catastrophe Insurance," *Journal of Financial Economics* 65, 283-305.