Available online at www.sciencedirect.com
ScienceDirect
Journal of Banking & Finance 32 (2008) 69–85
Journal of BANKING & FINANCE
www.elsevier.com/locate/jbf

On the pricing of intermediated risks: Theory and application to catastrophe reinsurance ☆

Kenneth A. Froot [a,*], Paul G.J. O'Connell [b]

[a] Harvard Business School, Nat. Bureau of Economic Research, Soldiers Field, Boston, MA 02163, United States
[b] FDO Partners, 134, Mount Auburn Street, Cambridge, MA 02138, LLC, United States

Available online 12 September 2007

Abstract

We model the equilibrium price and quantity of risk transfer between firms and financial intermediaries. Value-maximizing firms have downward sloping demands to cede risk, while intermediaries, who assume risk, provide less-than-fully-elastic supply. We show that equilibrium required returns will be "high" in the presence of financing imperfections that make intermediary capital costly. Moreover, financing imperfections can give rise to intermediary market power, so that small changes in financial imperfections can give rise to large changes in price.

We develop tests of this alternative against the null that the supply of intermediary capital is perfectly elastic. We take the US catastrophe reinsurance market as an example, using detailed data from Guy Carpenter & Co., covering a large fraction of the catastrophe risks exchanged during 1970–94. Our results suggest that the price of reinsurance generally exceeds "fair" values, particularly in the aftermath of large events, that market power of reinsurers is not a complete explanation for such pricing, and that reinsurers' high costs of capital appear to play an important role.

JEL classification: G20

Keywords: Insurance; Reinsurance; Intermediation; Risk management; Cost of capital; Hurdle rates; Insurance underwriting; Reinsurance underwriting

1. Introduction

What drives the prices of intermediated risk transfers? If capital markets were perfect, risks would flow costlessly from corporate hedgers to investors, and required returns would be "fair" in the sense that they would be determined entirely by investor preferences. For example, in a perfect market, firms would pay the riskfree rate to cede risks that are independent of aggregate wealth. In such a world, there would be no need for financial intermediation. Intermediaries, whose job is to distribute, transform, and inventory risk, could add no value. And under perfect markets there would be no rationale for corporate hedging in the first place. As Modigliani-Miller argued, firms would be indifferent between ceding risk (e.g., hedging) and financing risk (e.g., raising equity) at fair prices. So, for example, firms would never cede risks that were independent of aggregate wealth at a rate greater than the riskfree rate.

In practice, of course, markets are far from perfect. These imperfections at once give rise to firms' desire to cede risk and intermediaries' ability to profitably assume risk. For example, investors may be at a competitive disadvantage when it comes to evaluating and monitoring risks that are non-standardized and informationally opaque. If

☆ An earlier version of this paper was prepared for the NBER Conference on The Financing of Property/Casualty Risks, Palm Beach FL, November 21–23, 1996. We are grateful to Peter Diamond, Chris Milton, Julio Rotemberg, Jeremy Stein, Rene Stulz and conference participants for helpful comments, to David Govrin, Chris McGhee, Denis Medvedsek, Brian Murphy, and Barney Schauble for help in obtaining and understanding reinsurance data, and to the NBER, the Global Financial Systems Project, and Division of Research at Harvard Business School for generous research support. All errors are our own.
* Corresponding author. Tel.: +1 617 495 6677; fax: +1 617 496 7357.
E-mail addresses: kfroot@hbs.edu (K.A. Froot), poconnel@fdopartners.com (P.G.J. O'Connell).

0378-4266/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2007.09.008

forced to finance such risks directly, investors would charge a high rate. A cheaper solution might be for intermediaries to warehouse these risks, with investors financing the intermediaries. Intermediaries can do this and still add value because they provide evaluation and monitoring services. However, although intermediaries may reduce deadweight financing costs, they are unlikely to eliminate them entirely in their own financing needs. Lack of standardization and opacity will continue to be present. As a result, financial intermediation may occur, but the required return on non-standard and opaque intermediated risks will be high. Moreover, intermediaries' capacity for these risks will be less-than-perfectly elastic. In other words, the required return on non-standard and opaque risks that are independent of aggregate wealth will be greater than the riskfree rate and intermediaries will require successively greater returns for bearing additional quantities of such risk.

The industrial structure of intermediation may also be affected by financial imperfections. Bigger intermediaries may conserve on costly external finance because they are better able to diversify risks and fund investment opportunities of a given size. If so, then financing imperfections become a source of increasing returns to scale for intermediaries. Although small financing imperfections can generate only small returns to scale, they can nevertheless generate large increases in market power. The implication is that, under imperfect competition, even small financing imperfections can have large impacts on the equilibrium price of intermediated risk.

In this paper, we model the equilibrium pricing of risks that are non-standardized and opaque. In our view, firms wish to cede risks to economize on financing/investment costs. Because intermediaries specialize in bearing these risks, they can assume them at lower cost than investors, albeit at higher cost than "fair" value. The higher required returns paid by firms ceding these risks are a result of the costs intermediaries bear in funding themselves and the barriers to entry created by the financial imperfections intermediaries face.[1]

Based on these ideas, the model derives a firm's downward-sloping demand for hedging. This demand for hedging is a function of the financing imperfections facing the firm, the amount of financial slack the firm has initially, and the volatility of the risks facing the firm. The model is then used to derive an intermediary's upward-sloping supply of hedging capacity. The intersection of demand and supply is the equilibrium transaction price of intermediated risk. We show that the financing imperfections make the required return on this risk high. It is also clear how the risk profile and financial slack of firms and intermediaries affect conditions equilibrium price and quantity. Finally, we demonstrate how market power of intermediaries can interact with firms' and intermediaries' financing imperfections to raise the cost of hedging intermediated risk even further.

To motivate empirically our model of these issues, we examine one particular market for intermediated risk – that of catastrophe reinsurance. In this market, insurers purchase reinsurance contracts from reinsurers. Under these contracts, reinsurers agree to pay insurer damages resulting from natural perils such as hurricanes and earthquakes. Reinsurers pool these risks in and across their portfolios, but are unable to diversify them fully. This is because potential cat losses are large relative to reinsurer capital.[2] Given the magnitude of potential cat losses, one would expect insurers and reinsurers to hedge cat risk by finding investors with whom to share it. Yet, in fact, insurers and reinsurers tend to retain cat risks. Perhaps because catastrophe risks are neither standardized nor transparent, investors have historically been unwilling to share them directly. As a result, these risks yield high returns and are financed exclusively by insurers and reinsurers – both intermediaries who must find their own costly financing. In other words, the market for catastrophe reinsurance is an intermediated market in which the required return appears high, yet little direct risk transfer to investors occurs.

The market for catastrophe risk is particularly well suited to our analysis because catastrophe exposures are (arguably) independent of the risks on financial assets and because they can be measured using objective scientific models. If cat risks are diversifiable with respect to aggregate wealth, their "fair" required excess return is equal equivalent to the rate of actuarially expected loss. In other words, the total return for bearing diversifiable cat risk exposure should be the riskfree rate. Furthermore, quantitative and objective modeling of the probabilities of catastrophic losses is possible. This means that we can actually calculate the "fair" price of catastrophe reinsurance contracts, and use this to benchmark observed transaction prices. This is the exercise we undertake in the empirical section of this paper. Our benchmark prices come from an extensive set of reinsurance contract data from Guy Carpenter & Co., the largest broker of catastrophe reinsurance worldwide. These transactions data cover a significant fraction of the US catastrophe reinsurance market over the period 1970–1994 and allow us to explore the properties of equilibrium prices and quantities of cat risk transfer.

To preview our empirical findings, the average premiums (i.e., prices) on catastrophe reinsurance are considerably above our estimate of actuarial value. Cat risk, therefore, yields an expected return well in excess of the riskfree rate. Furthermore, we show that prices and quantities are negatively correlated. Both facts suggest that the

[1] Investors more readily bear standardized, transparent exposures, such as major currencies or stock indexes. This reduces the marginal cost of intermediation and the resulting potential for intermediary market power. Consequently, the supply of intermediary capacity will be highly elastic with respect to such risks.

[2] A single catastrophic event (such as a large hurricane or damaging earthquake) can generate potential insured losses of up to $100 billion in the US. Estimates of total capital and surplus of all US insurers is approximately $239 billion; the capital for reinsurers worldwide is estimated at $57 billion. See Froot (1999) for an overview of the market for catastrophe reinsurance.

supply of capital to bear cat risk is not perfectly elastic. Naturally, these results depend on our ability to model accurately the ex ante distribution of catastrophe losses. However, even if our empirical approximation of this distribution is flawed, the results from our analysis agree with the distribution of losses that comes out of commercially available models. These more sophisticated models, produced by specialized cat-modeling firms, suggest that current prices of catastrophe reinsurance are well in excess actuarial loss estimates.

Indeed, prices are particularly high in the aftermath of major catastrophic events. We find evidence that such price increases are driven by decreases in the supply of reinsurance as well as increases in demand.[3] The price impact of a shift outward in demand could be attributable to a combination of enhanced reinsurer market power and increased reinsurers' cost of capital. But part of the price impact is also driven by a shift backward in supply, suggesting that increases in reinsurers' costs of capital must be an important element.

We estimate the elasticity of demand to be between -0.2 and -0.3, while the elasticity of supply is approximately 7. The latter figure suggests that a 1% increase in premium above actuarially expected losses leads to a 7% increase in reinsurance capacity supplied. This response seems small relative to what capital markets could be expected to provide for liquid instruments. For example, a similar 10% increase in the interest rate on several billion dollars of one year risky corporate debt (equivalent to an increase in yield of about 50 basis points when interest rates are 5%), with no change in riskless rates or credit quality would presumably increase the supply of investors by more than 70%.[4,5]

We also present evidence that changes in reinsurers' capital costs affect supply. Specifically, we find that the supply for a given contract is reduced (i.e., the reinsurance premium is increased) when: (i) variance of losses under that contract is greater; and (ii) the covariance of losses under that contract with the loss distribution of reinsurer portfolios is greater. (In both cases, we hold constant demand factors such as the variance of insurer exposures and recent insurer losses, as well as reinsurer losses). These effects would have no impact on supply if financial markets were perfect. Because reinsurer financing imperfections may promote market power, it is difficult to disentangle the relative importance of financing imperfections versus market power in these results.

[3] Doherty and Smith (1993) present evidence that insurance markets are less competitive when prices are high.

[4] Note that it is appropriate to hold riskless interest rates constant in this example to the extent that changes in catastrophe prices appear independent of other financial market returns (including interest rates). As one might guess, actual correlations are neither statistically nor economically different from zero. See Froot et al. (1995).

[5] Indeed, an adjustment for credit quality would, if anything, make this investment even more attractive. That is, after a catastrophic event, reinsurers may receive, say, 10% higher premiums even though their credit quality has probably declined.

The rest of the paper is structured as follows. Section 2 lays out the model of supply and demand in an intermediated market for risk. Section 3 describes our strategy for implementing this model. Section 4 describes the data. Sections 5 and 6 then present estimation methodology and results. Section 7 concludes.

2. A model of hedging demand and supply

In this section, we model the price and quantity of risk transfer in an intermediated market. Our basic rationale for corporate hedging demand is that of Froot et al. (1993): hedging increases firm value by reducing costly fluctuations in investment spending and external fund raising.[6]

Here, however, we model equilibrium risk transfer. The equilibrium is interesting because intermediaries have limited capital and face costs of adding more, as in Froot and Stein (1998). Intermediary costs of external finance would seem natural since intermediaries are themselves corporations, subject to the same kinds of frictions that make corporate hedging desirable in the first place. We show below how these financial imperfections interact with market power to magnify the increase in the equilibrium price of intermediated risk.

The model has two time periods: present and future. In the present period, insurers ("firms") and reinsurers ("intermediaries") make insurance and reinsurance ("hedging") decisions. In the future period, catastrophe losses and firms' and intermediaries' stock prices are realized. Firms and intermediaries also make present-period hedging decisions with an eye toward their overall goal – maximization of shareholder value.

2.1. Firms' demand for hedging

In the present period, the ith firm begins with an inherited level of net internal assets, $w_{i,0}$. These net assets are exposed to uncertain shocks (e.g., catastrophe losses) given by ε_i. For simplicity, we assume ε_i is normally distributed, $\varepsilon_i \sim N(1, \sigma_i^2)$.[7] The firm's exposure to this shock can be managed by hedging an amount q_i. We assume that the hedging contract is linear, so that the contract has a payoff of $q_i\varepsilon_i$.[8] Net assets in the future period are therefore

$$w_i = w_{i,0}[(1 - q_i)\varepsilon_i + q_i(1 - p_i)], \tag{1}$$

[6] A similar motivation for hedging can be found in Stulz (1984) and Diamond (1984) (which deals specifically with the role of diversification in reducing firm-wide risks).

[7] This distributional assumption is made for the sake of simplicity and has no effect on the basic results. Of course, normality is unlikely to be a good assumption for the distribution of catastrophe losses. In the empirical section below, we more accurately model the empirical distribution of catastrophe losses.

[8] This assumption does not affect the qualitative nature of the results but simplifies the analysis considerably. In practice, of course, most reinsurance contracts are excess-of-loss treaties (which are nonlinear in insurer losses – see Footnote 25 below). In the empirical section, we model the distribution of the nonlinear excess-of-loss contract payoffs.

where q_i can be interpreted as the hedge ratio, and p_i is the unit cost of the hedge contract *in excess of fair value.* Intuitively (and as we show below), $p_i = 0$ in a market with no costs of financial intermediation. In other words, fair value is defined as the price that would prevail in a perfect market with costless access to investors.

If risk management is to matter to a firm, the distribution of net internal assets across future outcomes of ε_i must affect stock prices today. To establish this linkage, we use the FSS formulation, which assumes that in the future, the firm has positive-NPV investment opportunities it wishes to protect.[9] The investment requires an expenditure of I_i (to be determined in the future period after w_i is realized). It provides a net return of $F(I_i) = f(I_i) - I_i$, where $f(.)$ is an increasing, concave function. Clearly, the funds for this investment must come from some combination of external sources, e_i, and internal sources, w_i, so that $I_i = e_i + w_i$. The problem for the firm is that external funds cannot be costlessly tapped – raising external funding generates convex costs, given by $C(e_i)$.[10]

If managers maximize firm value, then the value of the firm in the future is the solution to the investment/financing problem:

$$P(w_i) = \max_{I_i} Ft(I_i) - C(e_i), \quad \text{subject to } I_i = e_i + w_i. \tag{2}$$

FSS show that, under these conditions, $P(w_i)$ is an increasing concave function with $1 \leqslant P_w$ and $P_{ww} \leqslant 0$.[11] Intuitively, low levels of internal assets cause the firm to experience costly cuts in investment and/or costly attempts to raise external funding. If fluctuations in the value of internal net assets can be avoided through hedging, then the prospect of experiencing such costs is reduced. Thus, it is the concavity of the value function, $P(w_i)$, that makes risk management value-enhancing for the firm.[12]

From the perspective of the present period, the firm chooses its hedging policy so as to maximize expected future value of the company – $\max_{qi} V_i = E[P(w_i)]$, where the expectation is taken with respect to ε_i. (For simplicity, we ignore discounting.) The first-order condition for this problem defines the optimal amount of hedging, q_i^*:

$$\begin{aligned} E\left[P_w \frac{\mathrm{d}w_i}{\mathrm{d}q_i}\right] &= E[P_w(1 - \varepsilon_i - p_i)] \\ &= -\mathrm{cov}[P_w, \varepsilon_i] - pE[P_w] = 0. \end{aligned} \tag{3}$$

Using the assumption that ε_i is distributed normally, we can solve this equation explicitly for the hedge quantity demanded:[13]

$$q_i^* = 1 - \frac{p_i}{\theta_i}, \tag{4}$$

where $\theta_i \equiv w_{i,0} G_i \sigma_i^2$ represents the strength of demand – it is literally the insurer's marginal financing cost of retaining an additional unit of risk – and $G_i = G(w_i) = -E[P_{ww}]/E[P_w] \geqslant 0$ is effectively a firm-specific measure of risk aversion to fluctuations in w_i. It is easy to show that G_i is monotonically decreasing in w_i, $G_w \leqslant 0$, and that $G(\infty) = 0$. The better capitalized the firm, the lower its risk aversion and the less there is to be gained from hedging.[14]

Eq. (4) is our hedging demand equation. It shows that the optimal hedge ratio, q_i^*, is a decreasing function of price, p_i, an increasing function of the variability of the underlying exposures, σ_i^2, and a decreasing function of internal funds, w_i (through $G(.)$).

The term p_i/θ_i is the product of the expected excess return-to-variance ratio, p_i/σ_i^2 – essentially, the "alpha" on ε_i risk – multiplied by the level of firm risk tolerance, $1/w_{i,0}G_i$. This term can be interpreted as the firm's desired or "target" exposure to ε_i risk.[15] If $p_i > 0$, the firm should optimally retain some of its own exposures as a value-maximizing investment decision. If $p_i = 0$ the firm will optimally cede all its ε_i risk and hedge completely.

2.2. Intermediaries' supply of hedge capacity

As with a firm, an intermediary (i.e., reinsurer) begins with an inherited level of net internal assets, $w_{R,0}$. For simplicity, we assume that there is a single intermediary, but that this intermediary nonetheless prices competitively. If the intermediary exchanges risks with firms, its net assets will be exposed to a portfolio of ε_i's from cedents; from Eq. (1) above, the intermediary will assume from the ith cedent risk given by $q_{R,i}(\varepsilon_i + p_i - 1)$. The intermediary's future-period net internal assets are therefore given by:

---

[9] For an insurer, these investment opportunities might involve the competitive pricing of insurance policies to gain or protect market share, upfront funding of brokerage expenses, purchases of property, etc.

[10] FSS show how a convex cost function arises in the standard optimal contracting setting introduced by Townsend (1979) and Gale and Hellwig (1985). Other applications include Stein (1998) and Froot and Stein (1998). FSS also provide arguments as to why corporate agency and information problems result in the kind of convex cost function of external finance assumed here.

[11] Our notation for derivatives is $P_w = \mathrm{d}P/\mathrm{d}w$, and $P_{ww} = \mathrm{d}^2P/\mathrm{d}w^2$.

[12] In the FSS formulation, a firm that inherited a capital structure with lower leverage, would have greater internal net assets, all else equal. It would therefore be able to mitigate future costs of external finance. Thus low leverage would seem to be an inexpensive means of avoiding costly external finance. Froot and Stein (1998) remedy this imperfection by incorporating carry costs for net internal assets. These costs (which result from factors such as foregone interest tax shields and agency costs), make it expensive for the firm to solve its risk management problem through underleveraging.

[13] If x and y are normally distributed and $a(.)$ and $b(.)$ are differentiable functions, then $\mathrm{cov}(a(x), b(y)) = E_x[a_x]E_y[b_y]\mathrm{cov}(x,y)$. In the absence of normality, there is no convenient closed-form solution. Furthermore, the qualitative aspects of that hedge ratio will be the same as derived in the simple case of normality.

[14] Investors as a group represent the deep capital markets, and as such have elastic demand, so that for investors $G_i = 0$.

[15] That is, the unhedged fraction of exposure (i.e., retention) is $1 - q_i^* = p_i/\theta_i$. Froot and Perold (1996) define target exposure, and show that, in general, it is given by the ratio of excess return to total variance times the level of investor risk tolerance.

$$w_R = \sum_i q_{R,i}(\varepsilon_i - 1 + p_i) + 1, \tag{5}$$

where for simplicity we have normalized initial intermediary wealth to one, $w_{R,0} = 1$.

2.2.1. Supply with financing imperfections

In this subsection, we assume that intermediaries face the same sorts of financial-market imperfections that firms face. Thus, intermediaries have profitable internal uses of funds and costs of raising external funds, just as firms have. Intermediaries are also value maximizers, just as firms are. They therefore also solve the maximization problem given in Eq. (2), using w_R (5) in place of w_i (1).

The intermediary's present-period decision to supply hedge capacity is therefore the solution to the problem, $\max_{q_i} V_R = E[P(w_R)]$, where $P(w_R) = \max_{I_R} F(I_R) - C(e_R)$, subject to $I_R = e_R + w_R$. The resulting first-order condition of V_R for the ith exposure is just the negative of Eq. (3). Solving this we have

$$q^*_{R,i} = \frac{p_i - G_R \mathrm{cov}[\varepsilon_i, \varepsilon_{R,i}]}{G_R \sigma_i^2} = \frac{p_i}{G_R \sigma_i^2} - \beta_i, \tag{6}$$

where $G_R = G(w_R) = -E[P_{ww}]/E[P_w]$ where $w = w_R$ is the risk aversion of the intermediary, and where $\beta_i = \mathrm{cov}[\varepsilon_i, \varepsilon_{R,i}]/\sigma_i^2$ (with $\varepsilon_{R,i} = \Sigma_{j \neq i} q_{R,j}\varepsilon_j$) is the covariance between a unit of the ith risk and the weighted average of the other $K - 1$ risks in the intermediary's portfolio, where K is the number of insurers. Note that Eq. (6) indicates that the optimal capacity provided for the ith risk, $q^*_{R,i}$, depends on the capacities provided for all other risks, $q_{R,j\neq i}$. Following Froot and Stein (1998), we can solve the K equations represented by (6) for the $K \times 1$ vector of optimal supplies

$$q^*_R = \Omega^{-1} \frac{p}{G_R}, \tag{7}$$

where q^*_R is the $N \times 1$ vector of optimal supplies, Ω is the $K \times K$ covariance matrix of the ε_i shocks, and p is the $K \times 1$ vector of per unit prices. Eq. (7) just says that the optimal allocation of intermediary capacity is mean-variance efficient – increasing in return, and decreasing in covariance. Note that the optimal allocation for each $q^*_{R,i}$ depends on the entire vector of prices, p.

Rather than work with the full solution to the intermediary's portfolio problem in Eq. (7), however, we use the partial solution in Eq. (6). This latter condition is preferable since it does not impose full optimality, although it is consistent with it. Our interest here is not really whether intermediaries form mean-variance efficient portfolios, but whether internal funds, variance, and covariance importantly influence market prices.

Eq. (6) is the supply-curve analog of Eq. (4). It says that the optimal amount of exposure to a given risk is equal to the difference between the intermediary's "target" exposure, $p_i/G_R\sigma_i^2$, and the "pre-existing" exposure to that risk already in the portfolio, β_i. This latter term is just the coefficient in a regression of the portfolio return (excluding the ith exposure) on the ith risk factor. It therefore conveys how much exposure to the ith factor is contained in the preexisting portfolio. All else equal, the higher the preexisting exposure, the lower the willingness to supply additional capacity. If preexisting exposure equals target exposure, then it is optimal neither to assume – nor cede – any of the ith risk.[16]

Note that Eq. (6) says that if $G_R > 0$, intermediaries will supply positive capacity only if the ith risk is negatively correlated with the rest of the portfolio. Clearly, if intermediaries have plenty of internal funds, $G_R = 0$, and risks are priced at fair value, $p_i = 0$.

2.2.2. Supply with imperfect competition

Financing imperfections may have other, more indirect effects on the supply of intermediary capacity. Most importantly, these imperfections generate increasing returns to scale: the larger is intermediary size, the better it can conserve on costly external funds. This occurs because the investment opportunity set (given by $F(I_i) = f(I_i) - I_i$) remains constant when the size of the intermediary (represented by $w_{R,0}$) grows. Indivisibilities in the size of risky positions will also create increasing returns, because larger intermediaries will be better able to diversify their investments into many risky exposures. Both of these arguments suggest that market power will be increasing in the size of intermediary financing imperfections.

It is straightforward to extend the model to allow for market power. To do this, we employ a simple model of Cournot competition among N symmetric intermediaries. Each chooses an amount of capacity to provide to the ith firm, $q_{i,n}$ (where $\Sigma_n q_{i,n} = q_i$) to maximize value, $\max_{q,i,n} V_n = E[P(w_n)]$, given internal wealth in Eq. (5). The first-order condition for the nth intermediary is similar to (3):

$$E\left[P_w \frac{\mathrm{d}w_n}{\mathrm{d}q_i}\right] = \mathrm{cov}[P_w, \varepsilon_i] + \left(p_i + q_{i,n}\frac{\mathrm{d}p_i}{\mathrm{d}q_{i,n}}\right)E[P_w] = 0, \tag{8}$$

except that now we account for the infra-marginal decline in price resulting from an increase in quantity supplied, $q_{i,n}\mathrm{d}p_i/\mathrm{d}q_{i,n}$. Note that with such strategic intermediary behavior, there is no supply curve per se: the intermediary merely chooses the profit-maximizing place along each of the K demand curves, and charges accordingly.

2.3. Equilibrium

In the equilibria with financing imperfections and/or with imperfect competition, total demand will equal supply, so that

$$q^*_{i,R} = \sum_n q^*_{i,n} = q^*_i w_{i,0}. \tag{9}$$

However, the equilibrium prices and quantities will differ depending on the extent of financing imperfections and market power.

[16] See Froot and Perold (1996).

2.3.1. Equilibrium with financing imperfections

Using (9) along with Eqs. (4) and (6), the simplest solution for equilibrium price and quantity is

$$p_i^* = \gamma_i G_R \sigma_i^2,$$
$$q_i^* = 1 - \frac{\gamma_i G_R \sigma_i^2}{\theta_i}, \qquad (10)$$

where γ_i is the exposure of the entire intermediary portfolio to the ith risk factor, $\gamma_i = \mathrm{cov}[\varepsilon_i, \varepsilon_R]/\sigma_i^2$, and $\varepsilon_R = \Sigma_i q_{R,i}\varepsilon_i$. Note that this total exposure equals the preexisting exposure to i, plus the size of the position in the ith risk, $\gamma_i = \beta_i + q_{R,i}$. For given β_i, equilibrium prices are increasing in the quantity of the ith exposure intermediary portfolios have to absorb. Here, because we have competitive behavior, price is equal to intermediary marginal cost, $p_i^* = \gamma_i G_R \sigma_i^2 \equiv mc_i$.

The results in Eq. (10) have intuitive properties. First, if intermediaries have effectively no exposure to the ith risk factor, then: $\gamma_i \equiv 0$; fair prices prevail, $p_i^* = 0$; and the ith firm hedges fully, $q_i^* = 1$. The ith firm effectively absorbs no intermediary capacity, so it is as though intermediation is costless. Also, high levels of intermediary capital imply that $G_R = 0$, again making intermediation effectively costless.

Second, if firms have ample internal funds, they face small costs of external finance, i.e., $G_i = 0$. In this case, demand is perfectly elastic at $p_i^* = 0$. Intermediaries will be able to supply capacity only to the extent γ_i is zero. Otherwise, no risk will be exchanged, $q_i^* = 0$, since it is cheaper for firms to retain exposures. To buffer their risks, firms would either prefund with large amounts of equity or issue equity contingent on catastrophes. In practice of course, either strategy will be costly: large equity buffers are tax-inefficient and promote agency problems and takeover pressures, while issuing equity in bad times is difficult and costly due to heightened informational asymmetries. This is just another way of saying that internal funds are scarce in practice.

Finally, the equilibrium has the property that as G_R and G_i converge to zero, price also goes to zero, but quantity becomes indeterminate. This is just the limiting case of Modigliani-Miller: if capital markets are perfect, the structure of financing does not matter.

2.3.2. Equilibrium with both imperfect competition and financing imperfections

Using the demand curve in Eq. (4), the equilibrium condition in (9), and the fact that the N intermediaries are symmetric, the solution to the imperfectly-competitive intermediary's problem in (8) is just:

$$p_i^* = \lambda\theta_i + (1-\lambda) mc_i,$$
$$q_i^* = (1-\lambda)\left(1 - \frac{mc_i}{\theta_i}\right), \qquad (11)$$

where $\lambda \equiv 1/N + 1$ is an increasing measure of market power, and $mc_i \equiv \gamma_i G_R \sigma_i^2$ is the (symmetric) marginal cost for the nth intermediary's investment in the ith firm's risk.

Eq. (11) is the standard solution to the Cournot problem with linear demand. Note that $\lambda = 0$ under perfect competition. In this case, the equilibrium in (11) converges to the perfectly competitive outcome in (10). With nonzero market power, $\lambda > 0$, price will be above intermediaries' marginal cost.

Eq. (11) is interesting in this context for several reasons. First, the greater is market power, λ, gives demand (given by θ_i) more scope to raise price and reduce quantity. Second, if market power is enhanced by greater intermediary financing imperfections, we have that $\lambda_G > 0$, where $\lambda \equiv \lambda(G_R)$. Third, by definition, demand and marginal costs are increasing functions of insurer and intermediary financing imperfections, respectively. That is, $\mathrm{d}\theta_i/\mathrm{d}G_i > 0$ since $\theta_i \equiv w_{i,0} G_i \sigma_i^2$, and $\mathrm{d}mc_i/\mathrm{d}G_R > 0$ since $mc_i \equiv \gamma_i G_R \sigma_i^2$. This implies that the impact on price of an increase in financing imperfections is

$$\frac{\mathrm{d}p_i}{\mathrm{d}G} = \lambda w_{i,0}\sigma_i^2 + (1-\lambda)\gamma_i\sigma_i^2 + (\theta_i - mc_i)\frac{\mathrm{d}\lambda}{\mathrm{d}G}. \qquad (12)$$

An increase in financing imperfections has three distinct effects, given by the terms in Eq. (12). The first term shows that the demand for hedging by firms increases with financial imperfections. This is a substitution effect – the costs of direct firm financing rise. The second term shows the component of price due to the decline in the supply of intermediary capacity as a result of more costly intermediary finance. The third term is the magnification effect of financing imperfections. That is, all else equal, prices are most sensitive to changes in imperfections when demand is relatively high (θ_i large) and when intermediary financing imperfections are small ($mc_i = 0$).

2.4. Discussion

The model is intended to motivate the empirical section below where we attempt to estimate the slope of the short-run supply curve for catastrophic risk taking. However, our analysis raises a number of important issues which deserve separate mention.

The first point – one that is directly relevant to the empirical work – is that we have modeled the short-run equilibrium only. Capacity and institutional structure are taken as constant, fixed factors. Positive shocks to demand may raise price in the short run, but over time, this price increase stimulates investment in intermediary capacity. Subsequent price declines are likely to follow. Loosely speaking, the long-run supply curve will be more elastic than the short-run supply curve because short-run marginal costs exceed long-run marginal costs. We provide loose evidence of these elasticities for catastrophe-risk intermediation below.

The second point is that our model is incomplete, in that it begs the question of why intermediation in sectors such as insurance is structured the (costly) way it is, even if we take as given the existence of financing imperfections. High equilibrium prices are partly a result of intermediaries failing to fully pool firm exposures. This lack of complete diversification is inefficient and expensive, so our model produces only a second-best outcome. Are there arrangements that might generate a constrained first-best outcome in which all risks are pooled?

One way to avoid costly intermediation is through an industry-wide merger of firms. This would be more efficient in that it would better spread firms' deadweight costs of external finance. However, standard monopoly arguments suggest that such a merger might generate enough market power among firms to reduce social welfare below where it is in our model. Besides (and partly as a result), such mergers are generally illegal.

Another way to avoid costly intermediation would be through the use of inter-firm exchanges of risk. Indeed, the unconstrained first-best outcome would have firms costlessly transferring their risks directly to investors. However, given the informational intensity and non-standardized nature of insurance risks, adverse-selection, moral hazard and agency problems make such an equilibrium potentially very costly. In the absence of standardized packages of risk, firms (or investors who may wish to fund these exposures directly) would want a monitor to evaluate risks and verify outcomes. The monitoring function may be most efficiently housed in a small number of organizations specializing in these activities. And, of course, this is the spirit of the multi-firm/single-intermediary structure of our model.

The third and most important point is that the source of high prices has important consequences for policy. If the story we tell is correct, then innovation in insurer/intermediary financing has very potent effects on the price of intermediated risk, as demonstrated in Eq. (12). In the insurance and reinsurance sectors, for example, use of catastrophe bonds and indexes of insurer losses might enhance standardization and transparency. This would reduce both insurer and reinsurer costs of capital and, simultaneously, cut reinsurer market power, further lowering the costs intermediated risk transfer.

If, on the other hand, market power is high for other reasons (e.g., barriers to entry due to reputation, etc.), there may be gains to encouraging competition among intermediaries. Note, however, that no matter how important these other sources of market power, they have little effect on equilibrium prices in the absence of financial imperfections. To see this, note that without financial imperfections, both θ_i and mc_i are zero in Eq. (11). The existence of low-cost substitutes to firm/intermediary hedging transactions limits intermediary market power and the distortion in prices.

Finally, we have ignored transactions costs as a potential explanation of high hedging costs. In some cases, such as reinsurance, transactions costs are likely to be important. However, transactions costs cannot easily explain the variation in prices over time nor high prices levels (transactions costs increase bid/ask spreads around fair value).

3. Testing the model

From Eq. (4) above, the quantity of reinsurance demanded is a decreasing function of price, increasing function of variance, and a decreasing function of internal funds. For purposes of estimation, we represent the demand for reinsurance by insurer i at time t by the simple quasi-log linear form

$$\ln(1+p_{i,t}) = \alpha_{1,i} + \alpha_{11}\ln(q_{i,t}) + \alpha_{12}\ln(\sigma^2_{i,t}) + \alpha_{13}w_{i,t} + v_{i,t}, \tag{13}$$

where $p_{i,t}$ and $q_{i,t}$ are, respectively, the price and quantity measures defined earlier, $\sigma^2_{i,t}$ is the per unit variance of the insurer's exposure, $w_{i,t}$ is the level of internal funds available to insurer i, and the α's are coefficients to be estimated. The coefficient $\alpha_{1,i}$ represents an insurer-specific fixed effect, attributable to unobserved differences in insurer-specific willingness to bear catastrophe exposure. The elasticity of demand is given by the inverse of coefficient α_{11}.

If insurer capital markets were perfect, so that $G_i = 0$, then risk management policies would be independent of the level of internal funds, $w_{i,t}$, implying $\alpha_{13} = 0$. Perfect markets would also imply that the variance of individual company exposures, $\sigma^2_{i,t}$, would have no influence on the demand for hedging, $\alpha_{12} = 0$. Furthermore with perfect financial markets, demand is perfectly elastic, so $\alpha_{11} = 0$.

On the supply side, Eq. (6) tells us that capacity supplied is an increasing function of variance and portfolio covariance relative to variance and a decreasing function of reinsurer internal funds. Linearized, this equation takes the form

$$\begin{aligned}\ln(1+p_{i,t}) = \alpha_2 + \alpha_{21}\ln(q_{i,t}) + \alpha_{22}\ln(\sigma^2_{R,i,t}) \\ + \alpha_{23}\beta_{i,t} + \alpha_{24}w_{R,t} + \eta_{i,t},\end{aligned} \tag{14}$$

where $q_{R,i,t}$ is the absolute quantity of reinsurance supplied to reinsurer i, $\sigma^2_{R,i,t}$ is the per-unit variance of the claims ceded by insurers to reinsurers, $\beta_{i,t} = \mathrm{cov}[\varepsilon_{i,t}, \varepsilon_{R,i,t}]/\sigma^2_{R,i,t}$ is the intermediary's preexisting portfolio covariance with respect to the ith firm's ceded risks, and $w_{R,t}$ is the industry-wide level of financial slack in the reinsurance sector. There are no insurer-specific intercepts in the supply function. Note that the unit variance of the ith exposure assumed by the reinsurer, $\sigma^2_{R,i,t}$, is not the same as the variance of the ith insurer's exposures, $\sigma^2_{i,t}$. This distinction is necessary because the contracts in our data transfer exposures which are nonlinear functions of underlying insurer portfolios. The covariances in $\beta_{i,t}$ capture the fact that, all else equal, an increase in the correlation of risks across insurers reduces effective reinsurer capacity.

The elasticity of supply in Eq. (14) is the inverse of α_{21}. If there were no financial-market imperfections impeding the flow of capital into catastrophic risk bearing, then the properties of intermediary portfolios would have no impact on supply, so that $\alpha_{23} = \alpha_{24} = 0$. Note, however that under this null hypothesis, α_{22} would be expected to be nonzero if intermediaries are imperfectly competitive. Only under the hypothesis that both financial markets and competition are perfect would supply be perfectly elastic, $\alpha_{21} = 0$. This implies that an incipient increase in $p_{i,t}$ above zero would result in an infinite amount of capacity becoming available.

To construct the variables for these regressions, we used data on the pricing and risk exposure of a panel of reinsurance contracts brokered between 1970 and 1994. In the sections that follow, we describe our basic data sources, our methodology for measuring risk exposure, and the precise construction of our regression variables.

4. Data

Our data is built up from four sources. The basic information on catastrophe reinsurance pricing is provided by Guy Carpenter & Co. Information on the regional market share of insurers is developed from A.M. Best data on insurance premiums written by company. Our estimates of catastrophe frequency and severity are based on Property Claims Services (PCS) data on US catastrophe losses since 1949. Finally, interest rate and CPI data are collected from Ibbotson and Associates and the IMF respectively.

4.1. Guy Carpenter catastrophe treaty data

Our basic data come from Guy Carpenter's proprietary database of catastrophe reinsurance contracts. Guy Carpenter & Co. is by far the largest US catastrophe reinsurance broker, with a market share of between 30% and 80% during our sample. The contracts brokered by Guy Carpenter cover a variety of natural perils, including earthquake, fire, hurricane, winter storm and windstorm.

From these data we extract transaction prices and quantities of "excess-of-loss" reinsurance contracts. Excess-of-loss contracts are defined by a deductible ("retention") and a maximum possible loss ("limit"). To understand how such contracts work, consider an insurer which purchases a layer of reinsurance covering \$100 million in cat losses "in excess of \$200 million". These terms imply that if the insurer's losses from a single catastrophic event during the contract year exceed \$200 million retention, the layer is triggered. The reinsurer pays the insurer the amount of any losses in excess of \$200 million, with the loss capped at a limit of \$100 million.[17] By purchasing this contract, the insurer cedes its exposure to single-event catastrophe losses in the \$200–\$300 million range. In return for assuming this exposure, the reinsurer receives a premium payment. If the insurer wishes to cede a broader band of exposure, it could purchase additional layers – \$100 million in excess of \$300 million, \$100 million in excess of \$400 million, and so on.[18]

We examine a total of 489 contracts brokered for 18 national and 19 regional insurers over the period 1970–1994.[19] These reinsurance contracts cover insurer losses sustained as the result of a single catastrophe events. The duration of coverage for each contract is one year. Data on contract inception date, retention, limit, losses, and premiums, company purchasing coverage, are employed. All of the contract inception dates are at the start of a quarter. Most contracts have a single mandatory reinstatement provision.[20]

4.2. A.M. Best market share data

To determine the catastrophe exposure of each contract, we must calculate the distribution of contract losses, a random variable for each contract. To do this, we assume that, within each region, each company's exposure is proportional to insurance industry exposure within the region. We therefore first determine a distribution for insurance industry losses for each region (by event type), and second multiply this aggregate distribution by an individual insurer's market share to determine the distribution of insurer-specific faced by that company. Using this information, we can calculate the company-specific distribution of losses under each contract.

Our estimates of insurer market shares are developed using data from A.M. Best on insurance premiums written by company, by line-of-business, by state, and by year. We reduce these multiline market shares to regional catastrophe market shares by applying a modified Kiln Formula, which assigns regional weights to premiums in each line of business based on exposure to catastrophes of that line in that region.[21] For example, depending on the region,

[17] To help guard against moral hazard, excess-of-loss reinsurance contract typically require coinsurance. In practice, this effectively means that the insurer provides 5–10% of the reinsurance itself.

[18] If losses are given by the random variable l, retentions by R, and the limit by Lim, then the excess-of-loss contract pays $\max(0, \min(Lim, l - R))$, where Lim and R are known at the time the contract is struck. This contract is equivalent to a call portfolio – the combination of a purchase of a call struck at R (with payoff linked to $\max(0, l - R)$), and the sale of a call struck at $Lim+R$ (with payoff linked to $\max(0, l - Lim - R)$).

[19] Seven very small regional insurers were dropped from the original Guy Carpenter & Co. data. In some of the computations below, we focus in on a smaller number of national reinsurers, for whom data are available in every year.

[20] The reinstatement provision stipulates that, conditional on an event which triggers losses on the contract, the limit is to mandatorily reinstated (one time only) by the reinsurer after payment of a reinstatement premium by the cedent. It appears that this provision has had only a modest effect on prices, and we ignore its effects. Conversations with brokers suggest that observed prices are approximately 10% *lower* than they would have been without the reinstatement premium. This seems surprising (forward contracts are usually priced at zero), but if anything leads us to underestimate what premia would be in the absence of reinstatement provision.

[21] This is a common industry practice. The specific weights used in our Kiln formula are from Guy Carpenter & Co.

anywhere between 50% and 95% of homeowners premiums are considered as funding catastrophe exposure. The five US regions used for insurer market shares are the Northeast, Southeast, Texas, the Midwest, and California.[22] We apply this market share data to all 489 reinsurance contracts selected from the Guy Carpenter & Co. treaty database.

4.3. Historic catastrophe loss data from Property Claims Services

As mentioned above, we need to determine the distribution of industry-wide losses to calculate the catastrophe exposure of each contract. To do this, we estimate the distributions of catastrophe frequency and severity using data from Property Claims Services (PCS). PCS has catalogued all catastrophe losses on an industry-wide basis since 1949 by type and US region. The PCS data are widely used as an industry standard.

Prior to estimating the parameters of the frequency and severity distributions, two adjustments are made to the PCS data. First, the losses are converted to 1994 dollars using the CPI. Second, they are modified to take into account shifts in the portfolios of property exposed to loss over the period. A key component of the latter adjustment is the demographic shift towards California, Florida, and Texas that has characterized recent decades. These two adjustments are carried out by Guy Carpenter & Co. Both adjustments are important. Indeed, the second adjustment implies that the same size event in *real* dollars causes damages which have grown on average by 5% per year over the sample period.

4.4. Interest rate and CPI data

For the purposes of calculating the net present value of payment flows, we use Ibbotson and Associate's index of the return on 30-day US Treasuries. This is collected monthly from 1970:1 to 1995:4. The US CPI is taken from the IMF's, *International Financial Statistics*. The frequency is monthly, from 1970:1 to 1995:3.

5. Calculation of exposure and price

5.1. Exposure

In this section, we describe our method of estimating the catastrophe exposure embodied in each excess-of-loss contract. The estimation is carried out in three stages. First, the frequency and severity of each type of event and region are estimated by maximum likelihood for particular families of distributions. Second, a simulated event history is generated by repeatedly drawing from the fitted frequency and severity distributions. Finally, the payouts under each contract in each year of event history are calculated. The mean of the distribution of these payouts is our estimate of the "quantity" of reinsurance, $q_{i,k,t}$, embedded in that particular contract.

5.1.1. The frequency and severity of catastrophes

The first step towards calculating contract exposure is to estimate the frequency and severity of catastrophes using the adjusted PCS loss data. Altogether there are over 1100 catastrophes recorded by PCS. These events are classified into 10 categories: earthquake, fire, flood, freeze, hail, hurricane, snowstorm, tornado, thunderstorm and windstorm.[23] Many of these events are relatively minor: only 557 have adjusted losses in excess of $15 million, and only 107 have losses in excess of $100 million. Four categories of losses are well-represented in the set of large losses: earthquake, fire, hurricane and windstorm.[24] As our primary interest is in exposure to large losses, we confine attention to these types. Examination of the data reveals that there is some heterogeneity in the losses that arise from windstorms. In particular, a number of the windstorms refer to winter storms ("Nor'easters") in New England. Accordingly, we split the windstorm category into two subcategories: winter storm, defined to be a windstorm in New England in either the first or fourth quarter, and windstorm, defined to be all other occurrences of a windstorm.[25]

Having defined these five categories of events, we need to make some assumption about regional effects before we can estimate frequency and severity distributions. The simplest assumption would be that, for each catastrophe type, event occurrences are drawn from a single nationwide frequency distribution while loss sizes are drawn from a single nationwide severity distribution. Given the relative paucity of loss information, this approach helps by pooling the available data. However, the assumption of equal regional distributions is likely to be incorrect. For instance, hurricanes are much less likely to occur in California than in Florida, and the majority of earthquakes occur in California.

[22] The regions are comprised as follows: Northeast – Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Pennsylvania, Rhode Island, Vermont; Southeast – Florida, Georgia, Mississippi, North Carolina, South Carolina, Virginia, West Virginia; Texas – Texas; Midwest – Illinois, Indiana, Kentucky, Missouri, Tennessee; California – California.

[23] PCS classifies many events into more than one category. For instance, winter storms in New England, which have on occasion caused substantial damage, are classified first as windstorms, and second as hail, freeze or snowstorm.

[24] During the 1949–1994 sample period, there were no floods, snowstorms or thunderstorms with losses in excess of $100 million. Only one freeze had losses in excess of $100 million, a $307 million freeze in Texas in 1989. Three hailstorms and three tornadoes did produce losses in excess of $100 million, but these are all dated prior to 1970, and so do not appear in our regression analysis below.

[25] The assumption that winter storms do not afflict the Midwest may seem strange. The reason is that our regional market share data is calculated for the Midwest using only five states: Illinois, Indiana, Kentucky, Missouri, Tennessee. The Dakotas, Michigan, Minnesota, Wisconsin and other characteristically Midwestern states are excluded.

Table 1
Frequency and severity assumptions by catastrophe type

Type	Description of PCS data	Assumptions		
		Regions	Frequency (# of regional distributions)	Severity (# of regional distributions)
Earthquake	10 events, all in CA. Frequency appears throughout year	CA	1: Uniform across quarters	1
Fire	19 events, 12 in CA, 2 in MW, 3 in NE, 2 in SE. Frequency higher in fourth quarter, and different for CA. Severity comparable across events.	NE, MW, CA	2: CA and NE/SE/MW/ TX. Both uniform across quarters	1
Hurricane	48 events, 26 in SE, 22 in NE and TX. Most in third quarter. More severe in Southeast	NE, SE, TX	8: SE (4 quarterly) and NE/ TX (4 quarterly)	2: Southeast, Northeast/ Texas
Winter storm	35 events, in NE in quarters 1 or 4	NE	1: uniform across quarters 1 and 4	1
Windstorm	352 events, all regions. Frequency differs across regions, but severity is comparable	NE, SE, TX, MW, CA	20: one for each region and quarter	1

Assumptions for catastrophe frequency and severity distributions, based on catastrophe experience 1949–1994. A catastrophe is defined as an event that gives rise to \$15 million or more in insured losses. Column 2 gives a description of catastrophe occurrence by type, 1949–1994. NE denotes northeast, SE southeast, TX Texas, MW Midwest and CA California. Columns 3, 4, and 5 give the assumptions concerning the frequency and severity distributions. The number in the frequency and severity columns represents the number of separately-estimated distributions for that type. For example, the number "1" implies that all regions are pooled, and that a single, nationwide distribution is estimated.

As a result, we make specific assumptions regarding frequency and severity on the basis of a careful examination of the 1949–1994 catastrophe data. These assumptions are summarized in Table 1. A catastrophe is defined as an event that gives rise to \$15 million or more in insured losses. Column 2 summarizes the event history for each type. Column 3 reports the regions in which each event type is assumed to occur. Columns 4 and 5 indicate the number of regional frequency and severity distributions estimated for each type. Some of the constraints, such as the assumption that winter storms do not strike California, seem entirely reasonable. Others, such as the assumption that earthquakes do not strike outside California or that winter storms do not hit the Midwest, are less tenable (though see Footnote 25), and are dictated largely by data availability.

With the assumptions described in Table 1, there are 33 frequency distributions to estimate. We assume that the frequencies are Poisson distributed, and estimate the Poisson parameters by maximum likelihood (the estimates are equal to the mean number of events that occur per quarter). Table 2 presents the frequency results in four quarterly arrays, by type and region. The estimated frequencies accord with what one might expect. For example, hurricanes are most likely to occur in the third quarter.

Next we consider severity. There are six severity distributions, one for each of the catastrophe types identified in Table 1. We fit two alternative density functions to the empirical severity distribution of each type. The first is a lognormal distribution, with density function for losses l given by $f(l) = \exp\{-[\ln(l) - \mu]^2/2\sigma^2\}/[l\sigma\sqrt{(2\pi)}],\ l > 0$. The second is a Pareto distribution, with density function $f(l) = \alpha\beta^{a}/l^{(1+\alpha)},\ l > \beta$. Once again, the estimation is carried out by maximum likelihood. The fitted distributions are reported in Table 3. For earthquake, winter storm and windstorm events, the likelihood ratio test selects the Pareto distribution as the better fit, while for fire and hurricane events, the lognormal distribution is preferred. However, because the Pareto distribution tends to place a large amount of probability in the right-hand tail of the distribution, it does not perform well in attaching reasonable probabilities to large losses. For example, using the estimated Pareto density, the probability that a hurricane in the Southeast generates \$15 billion in losses (given that a hurricane occurs) is almost 10%, which appears somewhat high.[26] It might be preferable, therefore, to use the lognormal fit as the baseline severity distribution for all event types. This is the strategy we adopt.

5.1.2. Simulated event history

Using these frequency and severity distributions, we are able to simulate an "event history" of catastrophes. From this event history the distribution of payments under each excess-of-loss contract can be obtained.

If all events were drawn from the same distribution, the distribution of aggregate losses could be estimated parametrically, so that payments under any contract could be analytically derived. However, with different event types, each with a different even distribution, this approach One complicating aspect of the simulation is that a contract's payment is triggered by only a single event, even though that event could be one of five different peril types. The single-event clause is in effect a knockout provision, allowing the contract to mature following the first event that generates losses in excess of the retention. For example, it may be that earthquakes are the major large risk for a contract to trigger, but a large freeze in the Northeast in early Jan-

[26] Using PCS data, Cummins et al. (1999) argue that the Pareto distribution tends to overestimate the probability in the tail of catastrophe severity distributions, and that the lognormal fit is to be preferred on these grounds.

Table 2
Frequency of catastrophes, measured by their Poisson parameters, by quarter, type and region, 1949–1994

	NE	SE	TX	MW	CA	NE	SE	TX	MW	CA
	January–March					April–June				
Earthquake					0.054					0.054
Fire	0.031	0.031	0.031	0.031	0.125	0.031	0.031	0.031	0.031	0.125
Hurricane (SE)		0.000					0.043			
Hurricane (NE/TX)	0.000		0.000			0.033		0.033		
Winter storm	0.380									
Windstorm		0.652	0.326	0.500	0.304	0.196	0.457	1.109	0.935	0.000
	July–September					October–December				
Earthquake					0.054					0.054
Fire	0.031	0.031	0.031	0.031	0.125	0.031	0.031	0.031	0.031	0.125
Hurricane (SE)		0.370					0.130			
Hurricane (NE/TX)	0.283		0.283			0.033		0.033		
Winter storm						0.380				
Windstorm	0.174	0.065	0.152	0.326	0.000		0.283	0.326	0.370	0.130

Poisson parameter is equivalent to the mean number of catastrophe occurrences per quarter by type and region. If the frequency of each catastrophe type in each region is Poisson distributed – $f(n) = e^{-\lambda}\lambda^n/n!$, where n is the number of events that occur – then the numbers in the table are the maximum likelihood estimates of λ. NE denotes Northeast, SE Southeast, TX Texas, MW Midwest and CA California. Blank elements of the arrays are 0 by assumption (see Table 1).

Table 3
Fitted severity distributions by catastrophe type, 1949–1994

Distribution	Parameter	Earth quake	Fire	Hurricane (SE)	Hurricane (NE/TX)	Winter storm	Wind storm
n		10	19	26	22	35	352
Lognormal	μ	−2.100	−2.350	−1.233	−1.454	−2.440	−3.039
	σ	1.964	1.196	1.610	1.454	1.166	0.859
	Mean log–L	0.006	0.752	−0.662	−0.340	0.867	1.772
	$\Pr(l > \$5\text{bn})\%$	2.915	0.046	3.870	1.760	0.025	0.000
	$\Pr(l > \$15\text{bn})\%$	0.684	0.001	0.718	0.211	0.000	0.000
Pareto	α	0.476	0.541	0.337	0.364	0.568	0.862
	β	0.015	0.015	0.015	0.015	0.015	0.015
	Mean log–L	0.358	0.735	−0.854	−0.556	0.875	1.891
	$\Pr(l > \$5\text{bn})\%$	6.288	4.327	14.110	12.057	3.684	0.670
	$\Pr(l > \$15\text{bn})\%$	3.727	2.389	9.743	8.082	1.973	0.260

Results from fitting of lognormal and Pareto distributions to PCS event losses. PCS losses have been adjusted for inflation and population movements by Guy Carpenter & Co. A catastrophe event is defined giving rise to insured losses in excess of \$15 million. The density function for the lognormal is $f(l) = \exp\{-[\ln(l) - \mu]^2/2\sigma^2\}/[l\sigma\sqrt{(2\pi)}]$, $l > 0$, while the density function for the Pareto is $f(l) = \alpha\beta^{\alpha}/l^{(1+\alpha)}$, $l > \beta$. The parameters μ, σ, and α (not β, which is a fixed scale parameter set equal to \$15,000,000) are estimated by maximum likelihood. For a given catastrophe type, estimated mean log-likelihoods for the two distributions are comparable, and provide a means for choosing between them. The table also shows the probability that an event produces insured losses in excess of \$5 billion and \$15 billion respectively.

uary could trigger the contract, thereby knocking out the earthquake risk for the remainder of the year.

This knockout provision gives the contract a payment distribution that is very different from that which would apply if the contracts were instead written to cover aggregate losses (i.e., the sum of losses across events). It can also give rise to some paradoxical effects. For example, an increase in the frequency of winter storms may actually *reduce* the total exposure embodied in a single-event contract, since it may increase the probability that it matures following a winter storm rather than a devastating hurricane.[27]

We simulate a 1250-year event history. For each quarter, the following steps are followed.

1. The number of events of each type that occur in each region is randomly drawn from the relevant Poisson frequency distribution (Table 2).
2. For each event that occurs, a loss amount is randomly drawn from the relevant severity distribution (Table 3).
3. All the events that occur in the quarter are randomly sequenced in time.

The random sequencing of the events throughout the quarter is an approximation, at best. It is likely, for example, that winter storms occur more frequently in January than March. While it would be preferable to sequence the events on a time scale finer than quarterly,

[27] An appendix examining the value of the knockout provision is available from the authors on request.

too few events that have occurred since 1949 to allow estimation of this.

5.1.3. Contract exposure

The exposure of each excess-of-loss contract in our data can be calculated by examining its loss experience in each year of the simulated event history. To take an example, suppose we are considering a contract purchased by a national insurer with an April 1 inception date. Let L and R be the contract's limit and retention, and let $m_{i,k}$, $k \in \{\text{NE SE TX MW CA}\}$ be the ith insurer's market share in each of the five regions. The contract's exposure is measured as follows:

1. Split the event history into 1249 year-long periods measured from April 1 to March 31.
2. Consider each period in turn. If no event occurs in a period, move to the next period. Otherwise consider each event in sequence.
   (a) Let the first event be in region k, and let insured losses from this event be l.
   (b) If $m_{i,k}l > R$, the contract is triggered. Measure the reinsurance payment for this period as min $(L, m_{i,k}l - R)$, and move on to the next period. The contract is no longer in force.
   (c) If $m_{i,k}l < R$, no payment takes place, and the contract remains in force. Move on to the next event, or the next period if there are no more events.

This algorithm generates 1249 observations on the distribution of payments under the contract. The first moment of this distribution is the expected exposure to catastrophe losses. It is easy to derive various conditional loss distributions from the unconditional distribution, such as the distribution of hurricane losses, or the distribution of losses from events in the Northeast.

We label the expectation of the unconditional distribution $q_{i,t}$, the exposure embodied in company i's contract at time t. Thus, $q_{i,t}$ is the actuarially expected loss covered by contract i. We use $q_{i,t}$ to represent the quantity of reinsurance purchased.

5.2. Other variables

To calculate contract price, we begin with the premium paid for each contract. This is simply measured as the sum of the premiums paid for each layer. Typically, the premiums are paid on a quarterly basis over the duration of the contract. We discount these premium flows back to the contract inception date using the three-month Treasury Bill rate.

Once the NPV of the premiums is calculated, it is converted to 1994 dollars using the CPI deflator. Our measure of price is the net present value (NPV) of premiums divided by contract exposure. Thus the price of company i's contract at time t is

$$p_{i,t} = \frac{\text{NPV(Premiums)}}{q_{i,t}} - 1 \tag{15}$$

Given our definition of quantity, the price of the contract is expressed as a unit increment to actuarially fair value.

$\sigma^2_{i,t}$ is the variance of underlying insurer portfolios. We calculate it using the simulated event history and the regional market share information for each insurer. Specifically, in each year of the event history, we estimate insurer i's losses by multiplying the simulated losses in each region by i's regional market shares. This generates 1249 observations on the distribution of insurer i's losses, from which we calculate first and second moments. $\sigma^2_{i,t}$ is the variance of the simulated distribution (in millions of 1994 dollars).

$w_{i,t}$ is the level of insurer internal funds. This is generally a difficult variable to measure. Even if one could accurately measure corporate net internal assets, their endogeneity makes them behave in ways that are difficult to interpret. For example, if a firm anticipates hard times, it may raise outside funds early, leading to the appearance that internal funds are plentiful, instead of scarce. Our solution is to use (the negative of) catastrophe losses by firm year as an instrumental variable for changes in net internal funds. This measure is particularly useful because cat losses are both strongly exogenous and correlated with changes in total internal assets. Unfortunately, we do not actually observe each insurer's catastrophe losses. Instead we infer their loss experience by combining *actual* catastrophe loss history, as measured by PCS, with insurer i's regional market shares. For each event, the loss amount recorded by PCS (in billions of 1994 dollars) is multiplied by the insurer's market share in the loss region to generate that insurer's losses. Internal funds $w_{i,t}$ are assumed to be depleted by the full amount of a loss for 8 quarters following the loss, after which time the impact of the loss on internal funds is zero.

$\sigma^2_{R,i,t}$ captures the variance of claims ceded by insurers to reinsurers. It is calculated in a manner similar to $\sigma^2_{i,t}$. The distribution of the payments under each contract is tabulated by examining the claims in each year of the simulated event history. As already discussed, $q_{i,t}$ is the first moment of this distribution. $\sigma^2_{R,t}$ is its variance, scaled by $q^2_{i,t}$.

Turning to $\beta_{i,t}$, its numerator is the covariance between per-unit claims under company i's contract, and the quantity-weighted sum of per-unit claims under all other outstanding contracts (i.e. with all insurers other than i). Its denominator is the variance of per-unit claims under company i's contract. The fact that $\beta_{i,t}$ is constructed using per-unit moments is important – it is intended to capture the covariation in contract returns, rather than total contract payouts. Accordingly, we scale payouts by their first moment prior to calculating $\beta_{i,t}$. Thus the numerator is obtained by calculating the covariance between contract i's payout and the sum of payouts on all contracts other than i from the simulated event history, and then dividing by $q_{i,t}\Sigma_j q_{j,t}$, $j \neq i$. The denominator is simply $\sigma^2_{R,i,t}$. We also consider a second measure of covariation, $\gamma_{i,t}$, which is calculated in a similar fashion to $\beta_{i,t}$, except that the covariance in the numerator is between contract i and a portfolio consisting of *all* contracts, including i.

Finally, $w_{R,t}$ is the level of internal funds available to the reinsurance industry. As with $w_{i,t}$, this is generally a difficult variable to measure. We use the total of reinsurance payments as reported by Guy Carpenter, scaled by Guy Carpenter's market share, as our measure of industry-wide reinsurance losses. Industry funds are assumed to be depleted by the full amount of any claims for 8 quarters following the claim, after which time the impact of the loss on industry funds is zero. The negative of this quantity, expressed in 1994 dollars, is $w_{R,t}$.

6. Estimation

6.1. Graphical analysis

It is useful to look first at the amount of catastrophe risk ceded by insurers during the sample period, and the average per-unit price at which this risk was ceded. Figs. 1 and 2 plot indices of industry price and quantity on a quarterly frequency from 1975:1 to 1993:4. The quantity series is the sum of the exposure ceded by four national insurers in each quarter for which data are available during the full sample, scaled by the total market share of these

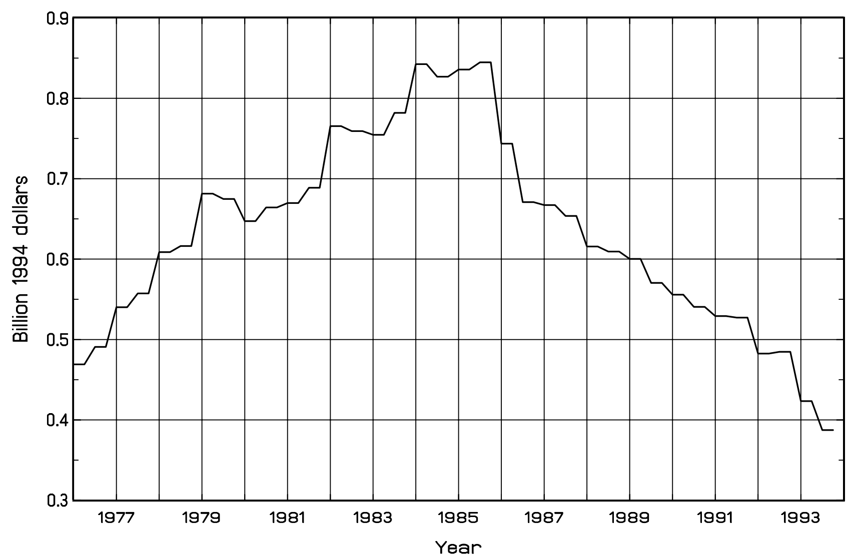

Fig. 1. Real quantity of catastrophe exposure ceded, 1975:1–1993:4.

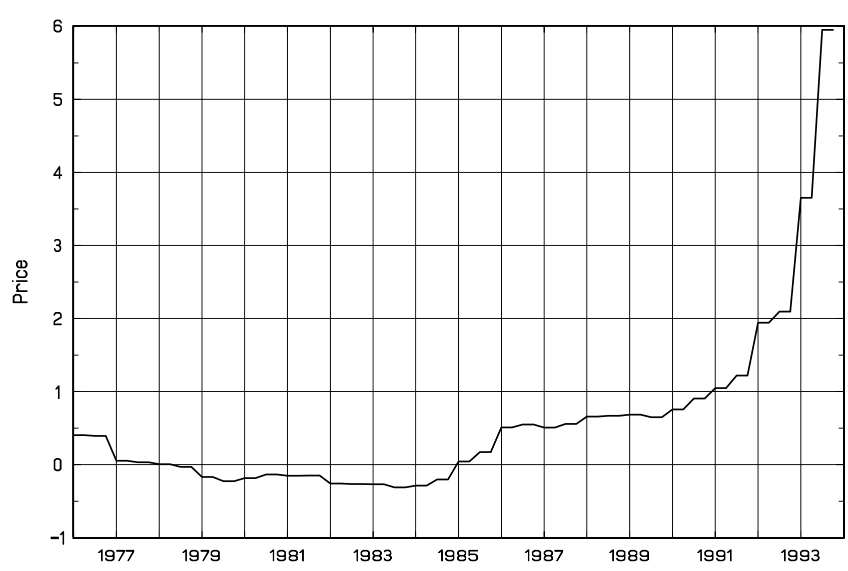

Fig. 2. Industry price per unit of ceded exposure, 1975:1–1993:4.

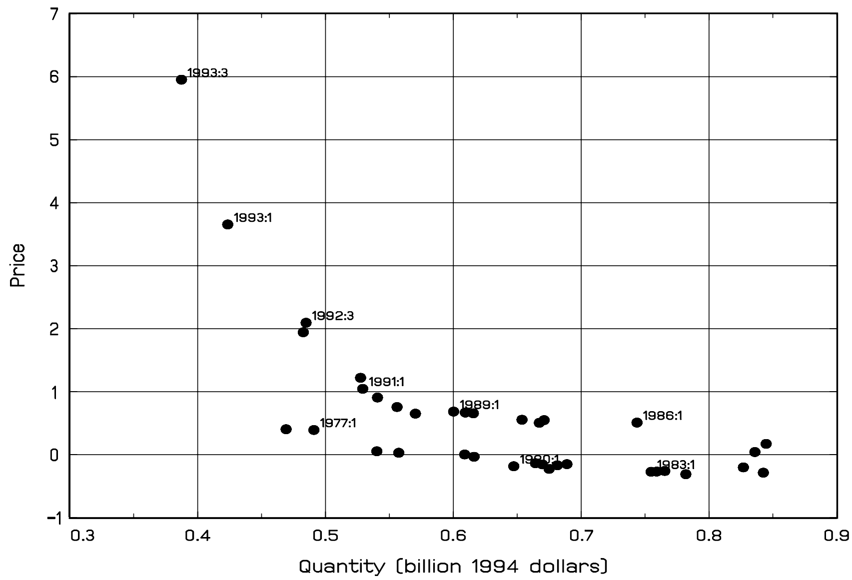

Fig. 3. Industry price-quantity pairs, 1975:1–1993:4.

four insurers.[28] The price series is the quantity-weighted average of the prices paid by these four insurers in each quarter. Fig. 3 shows a scatter plot of price and quantity.

A number of features of these figures are noteworthy. First, it appears that quantities rose and prices fell for much of the late 1970s and 1980s. Second, a startling rise in prices and decline in quantities took place beginning in the mid-1980s through the end of the sample. Indeed, in 1993, price was between 5 and 7 times its historical average. This will come as no surprise to industry observers. It is common to relate this price rise to the occurrence of a number of large events during this period, notably hurricane Andrew ($20 billion in losses) in August 1992, hurricane Hugo in 1989 and several windstorms in 1985–1986.

Fig. 4, which plots total catastrophe losses by quarter from 1970:1 to 1994:4 as measured by Property Claims Services, lends support to this view. In the period since 1994 (a period not covered by our data), the price of reinsurance has declined and quantity increased somewhat, notwithstanding the occurrence of the Northridge earthquake in January 1994.[29] From these observations, it is clear that there is considerable negative correlation between prices and quantities at frequencies of several years.

Hurricane Andrew is responsible for the largest catastrophe loss during our sample period. In light of this, it is of interest to look at the time series of prices around the time of this event. In particular, we can differentiate between the price-quantity reactions of those contracts heavily exposed to hurricane risk/Southeast risk and those with relatively less exposure. Table 4 contrasts the price and quantity responses. From Panel (a), we see that even those contracts with zero market share in the Southeast show large increases in price in the wake of Andrew. Panel (b) sorts contracts according to their hurricane exposure instead of by region. Contracts least exposed to hurricane losses that exhibit the largest increase in price. The results again suggest a negative correlation between prices and quantities.[30]

6.2. Estimation of supply and demand

Let $\mathbf{y}_{i,t}$ by the 2×1 vector of endogenous variables $[p_{i,t}\ q_{i,t}]$, and let $\mathbf{x}_{i,t}$ be the 5×1 vector of predetermined variables $[\sigma^2_{i,t} w_{i,t} \sigma^2_{Rit} w_{Rt} \beta_{i,t}]$. The structural Eqs. (13) and (14) above can be rewritten as

$$\mathbf{B}\mathbf{y}_{i,t} + \mathbf{\Gamma}\mathbf{x}_{i,t} = \boldsymbol{\alpha}_i + \mathbf{u}_{i,t}, \tag{16}$$

where

$$\mathbf{B} = \begin{bmatrix} 1 & -\alpha_{11} \\ 1 & -\alpha_{21} \end{bmatrix}$$

$$\mathbf{\Gamma} = \begin{bmatrix} -\alpha_{12} & -\alpha_{13} & 0 & 0 & 0 \\ 0 & 0 & -\alpha_{22} & -\alpha_{23} & -\alpha_{24} \end{bmatrix}$$

$$\boldsymbol{\alpha}_i = [\alpha_{1,i} \quad \alpha_2]$$

and $\mathbf{u}_{i,t}$ is a 2×1 vector of disturbances, distributed bivariate normally, with $E(\mathbf{u}_{i,t}\mathbf{u}'_{i,t}) = \mathbf{D}$, a diagonal matrix.

The reduced form of this system is obtained by premultiplying (16) by $\mathbf{B}^{-1}$, which yields

[28] The four insurers included in the industry indices purchased reinsurance through Guy Carpenter & Co. in each year from 1975:1 to 1993:4. They represent about 10% of the total market.

[29] Paragon Inc. produces a catastrophe price index shows the following prices since peaking in late 1994 at 2.47 (and beginning in 1/84 at 1.00): 1/1/95, 2.32; 7/1/95, 2.16; 1/1/96, 2.14; 7/1/96, 2.06.

[30] This negative correlation could be the result of our assumption that the distribution of losses is time invariant. Froot and O'Connell (1999) examine the hypothesis that the loss distribution may shift as a result of losses, but find little evidence that changing loss distributions explain the behavior of prices and quantities.

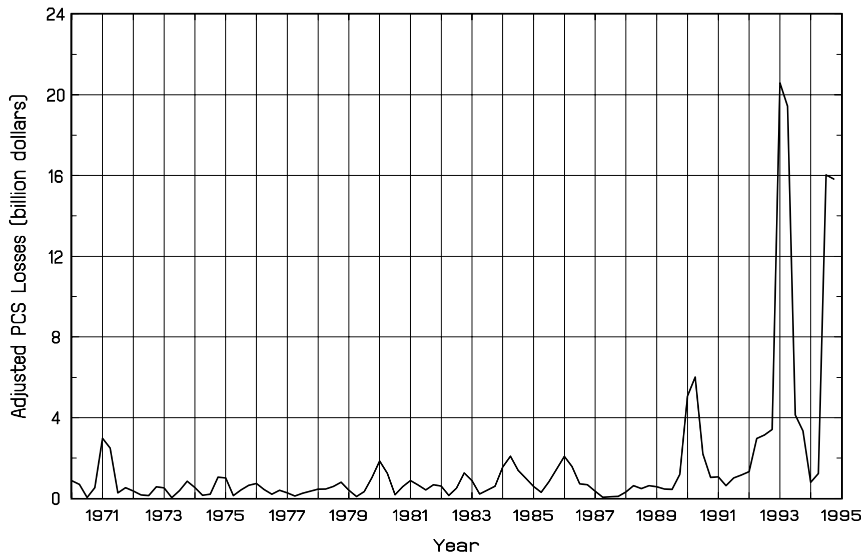

Fig. 4. Total adjusted PCS losses by quarter, 1970:1–1994:4.

$$\mathbf{y}_{i,t} = \mathbf{\Pi}'\mathbf{x}_{i,t} + p_i + \mathbf{v}_{i,t}, \tag{17}$$

where $\mathbf{\Pi}' = -\mathbf{B}^{-1}\mathbf{\Gamma}$ is a 2×5 matrix of reduced form parameters common to all companies, $\boldsymbol{\pi}_i = -\mathbf{B}^{-1}\boldsymbol{\alpha}_i$ is a 2×1 vector of reduced form intercepts for company i, and $\mathbf{v}_{i,t} = \mathbf{B}^{-1}\mathbf{u}_t$. We estimate this reduced form by full-information maximum likelihood (FIML). The conditional log-likelihood for company i is

$$L(\mathbf{B}, \mathbf{\Gamma}, \boldsymbol{\alpha}_i, \mathbf{D}) = -\frac{T}{2}\left(\ln(2\pi) - \ln|\mathbf{B}|^2 + \ln|\mathbf{D}| + \sum_{t=1}^{T}[\mathbf{B}\mathbf{y}_{i,t} + \mathbf{\Gamma}\mathbf{x}_{i,t} - a_i]'\mathbf{D}^{-1}[\mathbf{B}\mathbf{y}_{i,t} + \mathbf{\Gamma}\mathbf{x}_{i,t} - a_i]\right)$$

The log-likelihood for the full sample of 37 companies is therefore $L(\boldsymbol{B}, \mathbf{\Gamma}, \boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_{37}, \mathbf{D}) = \Sigma_i L_i$.

The FIML estimates are the values of $\boldsymbol{B}$, $\mathbf{\Gamma}$, $\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_{37}$ and $\mathbf{D}$ for which L is maximized.[31]

6.3. Results

Table 5 reports the FIML estimates of the structural coefficients, along with the estimated variances of the structural disturbance. The regressions are carried out both with and without company-specific intercepts. A likelihood ratio test easily rejects the common intercepts model, but nevertheless it is of interest to compare estimates across the two specifications. Standard errors for each coefficient are shown in parentheses.

Looking first at the demand specification, the elasticity of demand is estimated between -0.2 and -0.3, suggesting that, other things equal, a 1% increase in $p_{i,t}$ leads to a 0.25% reduction in quantity demanded. Lagged insurer losses exert an ambiguous effect on demand. The coefficient on $w_{i,t}$ is expected to be negative since lower internal funds implies higher reinsurance reservation prices. It is positive in the specifications without fixed effects and negative (the expected sign) when firm-specific fixed effects are included. To get a sense of magnitude, the coefficient in the first regression indicates that a reported loss of \$10 million by a company (any time over the preceding 8 quarters) decreases the price the insurer is willing to pay by 3.0%.

Increases in the variance of own-company exposure, captured by changes in σ^2_{it}, lead to significant increases in demand in all specifications. The point estimate in the first regression indicates that a 10% increase in the variance of a firm's risk exposure produces an increase in reservation price of 29.6%.[32] This number falls by a factor of two when company-specific intercepts are added in the latter specifications, yet it remains statistically significant. This is consistent with the hypothesis that financing imperfections play a role in the demand for catastrophe reinsurance.

Turning to supply, the elasticity is estimated to be on the order of 7 – ceteris paribus, a 1% increase in price produces a 7% point increase in quantity supplied. This

[31] We do not have a balanced panel of data, in the sense that at time t, $\mathbf{y}_{i,t}$ and $\mathbf{x}_{i,t}$ are observed over a (possibly empty) subset of the 37 sample companies. This does not present a problem for identification or estimation, as the simultaneity we are concerned with is within companies rather than across companies. The purpose of pooling the data is to obtain more efficient estimates of the structural parameters by imposing the constraint that they be equal (except for α_i) across companies. Note that we assume the disturbance terms are i.i.d. The rationale is that any contemporaneous and serial correlation in p_i and q_i ought to be captured in the predetermined variables.

[32] It should be noted, however, that there is no evidence of this effect in specifications 5 and 6, which are favored by the likelihood ratio test as the best characterizations of the data.

Table 4
Event study of hurricane Andrew

	(a) Southeast exposure			(b) Hurricane exposure		
	Mean exposure	Mean $\Delta\ln(p_{j,t})$	Mean $\Delta\ln(q_{j,t})$	Mean exposure	Mean $\Delta\ln(p_{j,t})$	Mean $\Delta\ln(q_{j,t})$
5 Most-exposed insurers	0.707	0.310	0.085	0.654	0.270	−0.030
5 Least-exposed insurers	0.000	0.334	−0.011	0.218	0.557	−0.138

Comparison of price responses in the year after hurricane Andrew (8/20/92–8/19/93) for different insurers. Panel (a) contrasts insurers which have high and low exposure to the Southeast (as measured by market share). Panel (b) contrasts insurers which have high and low exposure to hurricanes. The table shows the mean exposure and the mean price change of the 5 most extreme contracts in each case. The mean price change for the insurers with lesser exposure to the Southeast is calculated using all 14 of the insurers that have zero market share in that region.

Table 5
Estimates of structural model parameters

Demand	No fixed effects		Demand-equation fixed effects	
$\ln(q_{i,t})$	−4.856	−4.821	−3.420	−3.409
	(0.763)	(0.750)	(0.420)	(0.412)
$\ln(\sigma^2_{i,t})$	2.956	2.935	1.437	1.432
	(0.475)	(0.466)	(0.261)	(0.257)
$w_{i,t}$	3.017	2.978	−0.617	−0.624
	(1.728)	(1.710)	(0.809)	(0.801)
σ^2_v	5.769	5.689	1.197	1.190
	(1.807)	(1.762)	(0.293)	(0.287)
Supply				
$\ln(q_{i,t})$	0.130	0.134	0.144	0.145
	(0.019)	(0.018)	(0.020)	(0.018)
$\ln(\sigma^2_{R,i,t})$	0.679	0.682	0.692	0.693
	(0.030)	(0.030)	(0.030)	(0.030)
$\beta_{i,t}$	0.523		0.396	
	(0.261)		(0.264)	
$\gamma_{i,t}$		0.414		0.329
		(0.204)		(0.205)
$w_{R,t}$	−0.227	−0.226	−0.226	−0.224
	(0.020)	(0.020)	(0.020)	(0.020)
σ^2_η	0.214	0.214	0.215	0.215
	(0.014)	(0.014)	(0.014)	(0.014)
No. parameters	11	11	44	44
Mean log-L	0.502	0.503	0.950	0.950
No. observations	466	466	466	466

Estimated parameters of the demand and supply equations. See text for a complete description of the parameters and variables. Estimation is by full-information maximum likelihood (FIML). Standard errors in parentheses.

suggests that the marginal cost of reinsurer capital is upwardly sloped (though not strongly so). The coefficient on the variance of reinsurer-assumed exposure scaled by squared expected losses, $\sigma^2_{R,i,t}$, is positive and statistically significant. It indicates that, for a given assumed exposure, a 10% increase in the squared coefficient of variation increases prices by about 6.8%. In addition, the covariance term, $\beta_{i,t}$, is positive in all cases (and marginally statistically significant), so that exposures that are more correlated with reinsurance portfolios are priced higher. A 0.1 increase in the portfolio beta raises prices by 5.2%. Finally, reinsurer losses measured by $w_{R,t}$ enter negatively and significantly. A $100 million loss increases reinsurers reservation price by 2.3%. These coefficients seem of reasonable magnitude and are consistent with the capital market imperfections story.

7. Conclusions

We have traced out the implications of financing imperfections for equilibrium in markets for intermediated risks. Our results suggest that even small imperfections can lead to large deviations from fair pricing, particularly if these imperfections interact with intermediary market power.

In the case of catastrophe reinsurance, we used observed transactions to estimate reinsurer supply curves and insurer demand curves. These curves appear to shift in response to recent catastrophe losses, and to changes in the variance of insurer and reinsurer exposures. Furthermore, there is some evidence that reinsurers price both own variance and covariance of the risks they assume. This is consistent with the presence of capital-market imperfections. It does not rule out market power explanations, but suggests that market power cannot be the full explanation for high reinsurance prices.

References

Cummins, David J., Lewis, Christopher, Phillips, Richard D., 1999. Pricing excess of loss reinsurance contracts against catastrophic loss. In: Froot, K. (Ed.), The Financing of Property/Casualty Risk. University of Chicago Press.

Diamond, Douglas, 1984. Financial intermediation and delegated monitoring. Review of Economic Studies 51, 393–414.

Doherty, Neil A., Smith Jr., Clifford W., 1993. Corporate insurance strategy: The case of british petroleum. Journal of Applied Corporate Finance 6 (Fall).

Froot, Kenneth A., 1999. The limited financing of catastrophe risk: An overview. In: Froot, K. (Ed.), The Financing of Property/ Casualty Risk. University of Chicago Press.

Froot, Kenneth A., Murphy, Brian, Stern, Aaron, Usher, Steven, 1995. The Emerging Asset Class: Insurance Risk. Guy Carpenter & Co., Special Report, July.

Froot, Kenneth A., O'Connell, Paul G.J., 1999. The pricing of US catastrophe risk. In: Froot, K. (Ed.), The Financing of Catastrophe Risk. University of Chicago Press, pp. 195–232.

Froot, Kenneth A., Perold, André 1996. Determinants of Optimal Currency Hedging. Harvard Business School Working Paper no. 97-011.

Froot, Kenneth A., Scharfstein, David S., Stein, Jeremy C., 1993. Risk management: Coordinating corporate investment and financing activities. Journal of Finance 48, 1629–1658.

Froot, Kenneth A., Stein, Jeremy C., 1998. Risk management, capital budgeting and capital structure policy for financial institutions: An integrated approach. Journal of Financial Economics 47 (January), 55–82.

Gale, Douglas, Hellwig, Martin, 1985. Incentive-compatible debt contracts I: The one-period problem. Review of Economic Studies 52, 647–664.

Stein, Jeremy C., 1998. An adverse selection model of bank asset and liability management with implications for the transmission of monetary policy. Rand Journal of Economics 29 (3), 466–486.

Stulz, René, 1984. Optimal hedging policies. Journal of Financial and Quantitative Analysis 19, 127–140.

Townsend, Robert M., 1979. Optimal contract and competitive markets and costly state verification. Journal of Economic Theory 21, 265–293.