Bundling and Insurance of Independent Risks

Benjamin Davies and Richard Watt
University of Canterbury

Abstract

Risky prospects can often by disaggregated into several identifiable, smaller risks. In such cases, at least two modes of insurance are available: either (i) the disaggregated risks can be insured independently or (ii) the aggregate risk can be insured as one. We identify (ii) as risk bundling prior to insurance and (i) as separate, or unbundled, insurance. We investigate whether (i) or (ii) is preferable among consumers, insurers and the insurance market as a whole using numerical simulations. Our simulations reveal that separate contracts provide the socially optimal form of insurance when the insurer is able to charge the profit-maximising premia and has perfect information. Under asymmetric information with respect to consumers' risk aversion, we find that separation is again the dominant method of insurance in terms of the market share it represents.

Bundling and Insurance of Independent Risks

Benjamin Davies and Richard Watt
University of Canterbury

1 Introduction

Risky prospects can often be disaggregated into several identifiable, smaller risks. For example, a consumer may lose his home due to multiple random but imperfectly correlated events such as a flood, fire or earthquake. In such cases, at least two insurance options are available: either (i) the disaggregated risks can be insured independently or (ii) the aggregate risk can be insured as a whole. We identify option (ii) as "bundling" prior to insurance and (i) as "separate" or unbundled insurance. In this paper, we investigate the preference between bundling and separating among insurance demanders, suppliers and the insurance market as a whole.

"Bundling" as described is not a new issue. However, the concept has generally been taken to imply the sale of a package consisting of multiple separate products. Applied to insurance markets, standard bundling would involve simultaneous consideration of different risks under different insurance contracts provided by a single insurer, but where all of those contracts are sold together within a single bundle.[1] Bakos and Brynjolfsson (1999) show that bundling is efficient if there is some correlation between the consumer valuations of the bundled goods, which may imply that there is some efficiency in bundling together insurance products as protection against interdependent risks. Also related to insurance is the possibility that transaction costs can lead to a bundled set of contracts being priced lower than the sum of the individual components.[2]

Rather than investigate the packaging together of multiple insurance contracts, our goal in this paper is to examine the benefits of bundling risks together *before* insurance is sought and compare that to a scenario in which different risks are insured under different contracts simultaneously. Hence this paper interprets bundling as the insurance of multiple risks under a single contract. Thatcher and Clemons (2000) offer a simulation study in which multiple risks are bundled and then insured under a single contract in response to adverse selection in some of those risks. They show that such bundling leads to efficiency gains in the sense that the well-known negative effects of public disclosure of information on individual insurable risks can be overcome. We focus on risks that are independent so

[1]There is now an extensive literature on bundling in general that points to many different arguments for the benefits of selling several different products together. A recent and in-depth discussion of the bundling literature is offered by Sheikhzadeh and Elahi (2013).

[2]For a recent paper on the role of cost in bundling, see Evans and Salinger (2008).

as to remove the motive to purchase bundled contracts for the hedging or informational benefits.

Our central interest is in investigating why separating risks prior to insurance appears to be a much more prevalent feature of insurance markets than is bundling.[3] The structure of this investigation and of this paper as a whole is as follows. Section 2 outlines the decision problem faced by a consumer seeking to insure multiple risks through either a bundled or separate contracts. Section 3 describes the method by which we analyse the consumer's and insurer's preference between the two insurance options, and breaks down our simulations into three layered studies. Section 4 illustrates and interprets the results of our simulations. Finally, section 5 concludes with a discussion of our findings and identifies several opportunities for further research.

2 Analytical Framework

Assume that a consumer with initial wealth w_0 is exposed to two independent risks, both of which imply losses in wealth, that are identified by random variables L_1 and L_2. The consumer's aggregate risk is given by the sum $L = L_1 + L_2$. We assume that each loss L_i is a Bernoulli random variable with probability mass

$$\Pr(L_i = x) = \begin{cases} p_i & \text{if } x = l_i \\ 1 - p_i & \text{if } x = 0 \\ 0 & \text{otherwise,} \end{cases}$$

where $l_i > 0$ is a fixed loss size and $0 < p_i < 1$ for each $i \in \{0, 1\}$. We assume that $l_1 \leq l_2$ without loss of generality and that $l_1 + l_2 < w_0$. Since L_1 and L_2 are independent, the aggregate loss L partitions the world into four states: one in which no loss occurs, one in which only loss 1 occurs, one in which only loss 2 occurs and one in which both occur.

We consider the following two insurance options offered by a single insurer:

(i) The consumer buys separate contracts to insure each risk independently by choosing, for each $i \in \{1, 2\}$, a deductible minimum d_i above which all realisations of L_i are fully insured.

(ii) The consumer buys a bundled contract by choosing a deductible minimum D above which all realisations of L are fully insured.

[3]Koehl and Villeneuve (2001) remark that "in practice, insurance firms are specialized, for legal or strategic reasons, and even when they are not, they tend to offer contracts for different risks separately rather than real bundles." Indeed, browsing the product catalogue of several New Zealand-based insurance companies reveals few offerings of what we refer to as bundled contracts that protect against losses from mutliple sources.

The premium paid for each option is equal to the expected indemnity payout under that option multiplied by some insurer-set loading factor. Hence the premium for the bundled contract is given by

$$\pi_b(D) \equiv \lambda_b \, \mathrm{E}[\max\{0, L - D\}]$$

while the total premium payment for the two separate contracts is

$$\pi_s(d_1, d_2) \equiv \lambda_s(\mathrm{E}[\max\{0, L_1 - d_1\}] + \mathrm{E}[\max\{0, L_2 - d_2\}]),$$

where $\lambda_b \geq 1$ and $\lambda_s \geq 1$ are the associated loading factors. The consumer chooses the deductible minima D, d_1 and d_2 that maximise his expected utility under each option (i) and (ii). Thus his objective function under the bundled option is

$$\begin{aligned} \mathrm{E}[u(W_b)] &= \mathrm{E}[u(w_0 - \pi_b(D) - L + \max\{0, L - D\})] \\ &= \mathrm{E}[u(w_0 - \pi_b(D) - \min\{L, D\})], \end{aligned} \tag{1}$$

where u is his von Neumann-Morgenstern utility function and $W_b = W_b(D)$ is his wealth under option (ii). We assume that the consumer cannot purchase a negative amount of insurance nor receive a larger indemnity payment than his incurred loss. The optimal bundled deductible D^* therefore solves the constrained maximisation problem

$$\max_D \mathrm{E}[u(W_b)] \text{ subject to } 0 \leq D \leq l_1 + l_2. \tag{2}$$

The consumer's objective function under the separate option is

$$\begin{aligned} \mathrm{E}[u(W_s)] &= \mathrm{E}[u(w_0 - \pi_s(d_1, d_2) - L + \max\{0, L_1 - d_1\} + \max\{0, L_2 - d_2\})] \\ &= \mathrm{E}[u(w_0 - \pi_s(d_1, d_2) - \min\{L_1, d_1\} - \min\{L_2, d_2\})], \end{aligned} \tag{3}$$

where $W_s = W_s(d_1, d_2)$ is his random wealth under option (i). Again, we assume that each contract can provide neither a negative amount of insurance coverage nor more coverage than is required to indemnify the loss on which each contract is written. The optimal separate deductibles d_1^* and d_2^* therefore solve

$$\max_{d_1, d_2} \mathrm{E}[u(W_s)] \text{ subject to } 0 \leq d_1 \leq l_1 \text{ and } 0 \leq d_2 \leq l_2. \tag{4}$$

The consumer's insurance demand problem thus has two stages: first, he identifies the deductible minima D^*, d_1^* and d_2^* that maximise his expected utility under each option; second, he chooses the option that provides the highest level of expected utility at the relevant optima.

It is tempting to conclude immediately that the consumer must be better off insuring the two risks independently because doing so gives greater scope for choice than does the

Table 1: State probabilities and wealths from separate and bundled options

State	Probability	W_s	W_b		
			$0 \leq D \leq l_1$	$l_1 \leq D \leq l_2$	$l_2 \leq D \leq l_1 + l_2$
1	$(1-p_1)(1-p_2)$	$w_0 - \pi_s$	$w_0 - \pi_b$	$w_0 - \pi_b$	$w_0 - \pi_b$
2	$p_1(1-p_2)$	$w_0 - \pi_s - d_1$	$w_0 - \pi_b - D$	$w_0 - \pi_b - l_1$	$w_0 - \pi_b - l_1$
3	$(1-p_1)p_2$	$w_0 - \pi_s - d_2$	$w_0 - \pi_b - D$	$w_0 - \pi_b - D$	$w_0 - \pi_b - l_2$
4	$p_1 p_2$	$w_0 - \pi_s - d_1 - d_2$	$w_0 - \pi_b - D$	$w_0 - \pi_b - D$	$w_0 - \pi_b - D$

bundled option. However, it is not true that for each bundled deductible D there is a pair of separate deductibles (d_1, d_2) such that both options provide the same state-contingent wealth. This can be seen in Table 1, which shows the consumer's final wealths W_s and W_b in the four states described by the random variable L. If both options were to offer the same wealth in state 1 then we would have $\pi_b = \pi_s$. Similarly, if both options were to offer the same wealth in state 4 then we would have $D = d_1 + d_2$. But then the consumer's wealth in states 2 and 3 must be different. We conclude that the two options must have distinct state wealth profiles.

Given that neither insurance option clearly dominates the other in terms of state-contingent wealths, it is relevant to ask the following three questions.

1. Which option should consumers choose?

2. Which option is better for the insurer to offer?

3. What are the social welfare implications of such choices?

Real-world insurance markets exhibit both options. For example, consumers often buy home insurance that protects against fire, earthquakes and burglaries through a bundled contract. But insurers seldom bundle house insurance with, say, car insurance within a single contract. One explanation for this separation across risk categories might involve transaction costs. It is possible that insurance companies that primarily insure houses suffer greater costs in underwriting and auditing risks involving cars. We assume such costs away in order to concentrate on the fundamental characteristics of risk bearing as a possible explanation for bundling.

3 Method

In the first instance, we assume that the insurer sets the same proportional loading factor under both options (i) and (ii) described in the previous section. Under this assumption, any preference that the consumer may have for one option over the other cannot arise from a more affordable premium and must be in response to the stochastic structure of the associated indemnities. Identifying this preference is a matter of solving (2) and (4), and comparing the maximal objective function values. Unfortunately, the presence of the min function implies that (1) and (3) are not globally differentiable, and so precludes the use of calculus. We therefore restrict our analysis to numerical simulation. Such simulations allow us to clearly visualise the consumer's preference between bundling and separating by plotting his certainty equivalent wealth achieved under the two options.

Secondly, we extend these simulations to examine the case in which the insurer has perfect information with respect to the consumer's choice environment and is therefore able to set the loading factor under each option so as to maximise her expected profits. Identifying these loading factors analytically is infeasible because they depend upon the consumer's choice of deductible minima. In other words, the consumer's insurance demand problems are nested within the insurer's profit maximisation problems. Simulation is therefore a necessary but powerful tool for analysing the preference between bundling and separating when the insurer is able to maximise her expected profits because it does not require that we perform the analysis symbolically. Moreover, it allows us to visualise the insurer's preference between bundling and separating by plotting her associated level of expected profits under each option.

Finally, we relax the assumption of perfect information and consider an environment in which the insurer is imperfectly informed about the consumer's level of risk aversion.[4] Given the ubiquity of informational asymmetries in economies in general and insurance markets in particular, this third set of simulations offers the most realistic portrayal of how real-world insurance markets operate.

4 Simulation Results

In all that follows, we assume that the consumer has power utility

$$u(w) = \frac{w^{1-\gamma} - 1}{1 - \gamma},$$

where $\gamma > 0$ denotes his coefficient of relative risk aversion.[5] We also assume that he has initial wealth $w_0 = 20$, and that his endowed losses are identically distributed with $l_i = 5$

[4] Our assumption on asymmetric information is therefore different from that which is habitual in the literature, which revolves around risk characteristics rather than risk aversion. Making risk probabilities different over different individuals creates a far more complex problem, as there are two independent risks in the model for each insured individual. Even sticking with two-dimensional risks, there are up to four different categories of riskiness into which consumers may be classified. We prefer to stick with only two categories of insurance consumers, identified according to risk aversion types.

[5] We also simulated the case in which the consumer has exponential utility $u(w) = 1 - \exp(-aw)$ with coefficient of absolute risk aversion $a > 0$. The results of these simulations were identical to those with power utility, varying only in the values of γ and a that index the consumer's behaviour with respect to his level of risk aversion. The results from our simulations with exponential utility are available upon request.

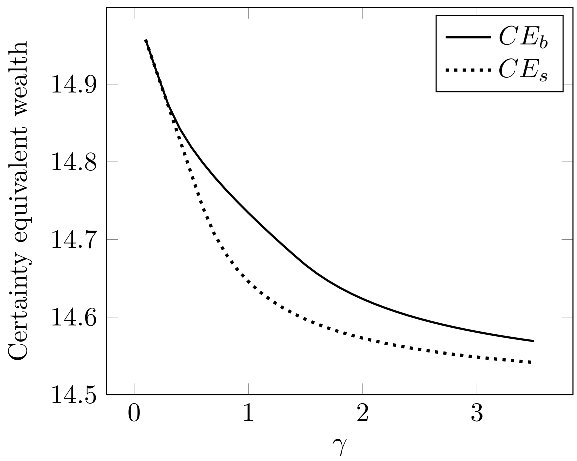

Figure 1: Consumer's certainty equivalent wealth with $\lambda_b = \lambda_s = 1.1$

and $p_i = 0.5$ for each $i \in \{1, 2\}$. These parameter values generate clear visual displays of the consumer's behaviour as will be shown throughout sections 4.1–4.3.[6]

4.1 Perfect Information with Common Loading Factors

Suppose first that the loading factor on each insurance option is fixed at $\lambda_b = \lambda_s = 1.1$. Numerical simulation allows us to examine how the consumer's preference between the two insurance options varies with respect to the value of γ. We consider such values over the interval $[0, 3.5]$. First, we compute the maximal levels of expected utility obtained under the two options by solving (2) and (4) numerically.[7] We then convert these levels into monetary units by computing the consumer's certainty equivalent levels of wealth

$$CE_b = u^{-1}(\mathrm{E}[u(W_b^*)])$$

and

$$CE_s = u^{-1}(\mathrm{E}[u(W_s^*)]),$$

where $W_b^* = W_b(D^*)$ and $W_s^* = W_s(d_1^*, d_2^*)$ are the state-contigent wealths achieved from the consumer's optimal contracts. Figure 1 plots these certainty equivalent wealths against the consumer's index of relative risk aversion γ. We observe that the consumer prefers the bundling option to separating for all values of $\gamma \in [0, 3.5]$ because the former provides a higher level of certainty equivalent wealth. This observation is a direct consequence of the following theorem.

[6]We also ran several simulations with asymmetric risks, changing either the probability or size of each risk. We did not find any fundamental differences in the final results as compared to the case in which the risks are identically distributed.

[7]The Python code used to generate the simulation data displayed in Figures 1–7 is available online at `https://github.com/bldavies/bundling-insurance`.

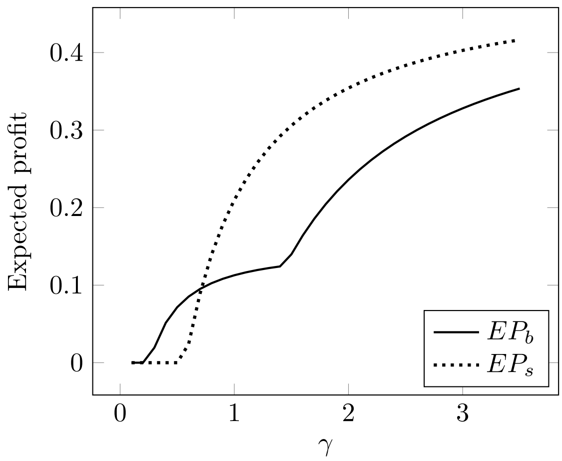

Figure 2: Insurer's expected profit with $\lambda_b = \lambda_s = 1.1$

Theorem (Arrow, 1963)**.** *If an insurance company is willing to offer any insurance policy against loss desired by the buyer at a premium which depends only on the policy's actuarial value, then the policy chosen by a risk-averting buyer will take the form of 100 per cent coverage above a deductible minimum.*

Since the bundled option describes exactly a contract that provides full coverage above a single deductible minimum, Arrow's theorem implies that it must be preferable to all other contract forms and thus, in particular, the separate option in which two deductibles are required.

We can also use the simulations to investigate the insurer's preference between the two insurance options. The insurer's expected profit under the bundled option is given by

$$EP_b = \pi_b(D^*) - \mathrm{E}[\max\{0, L - D^*\}];$$

that is, by her premium revenues less the expected cost of indemnifying the consumer's loss. Similarly, her expected profit under the separate option is given by

$$EP_s = \pi_s(d_1^*, d_2^*) - \mathrm{E}[\max\{0, L_1 - d_1^*\} + \max\{0, L_2 - d_2^*\}].$$

Figure 2 plots EP_b and EP_s against γ. We observe that the insurer's preference between bundling and separating changes depending on the magnitude of γ; if γ is small then the insurer prefers to supply the bundled option, while if γ is large then she prefers to sell separate contracts because doing so yields higher profits in expectation. Interestingly, the insurer's preference for separating is stronger than the consumer's preference for bundling. Figure 1 shows that the largest difference in the consumer's certainty equivalent wealths is about 0.09 units in favour of bundling at $\gamma = 1$, while Figure 2 shows that at $\gamma = 1$ the insurer earns about 0.1 units more profit from separating than bundling. Hence if $\gamma = 1$

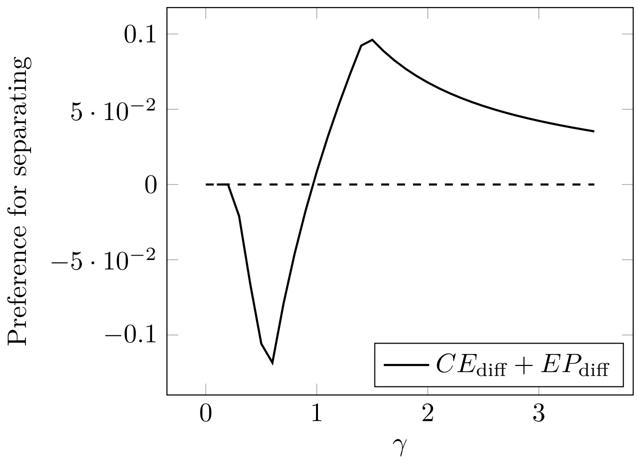

Figure 3: Social preference with $\lambda_b = \lambda_s = 1.1$

then trading the separate contracts provides the insurance market with about 0.01 units more surplus than bundling.

This social preference toward separating is investigated further in Figure 3, which plots the sum

$$CE_{\text{diff}} + EP_{\text{diff}} = (CE_s - CE_b) + (EP_s - EP_b) \tag{5}$$

against γ. Equation (5) provides a measure of the total market surplus that the separate option provides in excess of bundling. Figure 3 shows that this excess surplus is positive whenever $\gamma > 1$, thus implying that separating is socially optimal for all such γ. Indeed, if $\gamma > 1$ then the insurer earns more excess profits from separating than the consumer earns in excess certainty equivalent wealth from bundling. As a result, the insurer can afford to enact a Kaldor-Hicks transfer (e.g., by lowering λ_s) that makes the consumer prefer separating without forfeiting the excess profits earned from separating. That such a transfer exists suggests that the loading factors λ_b and λ_s play an important role in determining each party's preference between the two insurance options. We examine this role in more detail within the following section.

We note in passing that the social preference in favour of bundling at low values of γ is due simply to the fact that, for such γ, only bundled contracts are sold. This can be seen via Figure 2, in which the insurer earns zero profits from the separate option for $\gamma \leq 0.5$ because the consumer is insufficiently risk averse to buy the pair of separate contracts. The insurer's and social preference toward separating quickly becomes dominant as the consumer becomes more risk averse and demands a higher level of coverage under the separate option.

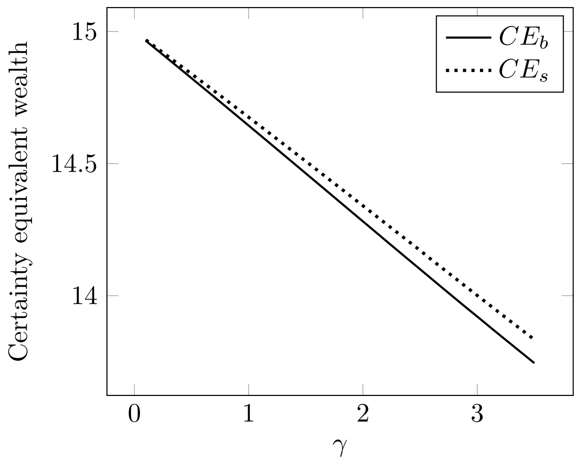

Figure 4: Consumer's certainty equivalent wealth with $\lambda_b = \lambda_b^*$ and $\lambda_s = \lambda_s^*$

4.2 Perfect Information with Profit-Maximising Loading Factors

Now suppose that the insurer is free to set the loading factors λ_b and λ_s that maximise her expected profits from the two insurance options. If the insurer has perfect information about the consumer's decision environment then she can compute the profit-maximising loading factors λ_b^* and λ_s^*, subject to the demand schedules implied by the solutions to (2) and (4). Both λ_b^* and λ_s^* can be found numerically, and we can use the numerical solutions to investigate the consumer's and insurer's preference between bundling and separating.

Figure 4 shows the consumer's certainty equivalent wealths from bundling and separating when the insurer is able to set the profit-maximising loading factors. In contrast to when $\lambda_b = \lambda_s = 1.1$, we see that the consumer strictly prefers the separate option for all values of γ when $\lambda_b = \lambda_b^*$ and $\lambda_s = \lambda_s^*$. This change in preference between the two cases shows that Arrow's theorem is sensitive to the assumption that all contracts have the same loading factor. The reason why we observe so many separate contracts in real-world insurance markets despite Arrow's theorem may simply be that the theorem ignores the role of the insurer in determining the equilibrium insurance contracts.

Separating is also strictly preferred by the insurer when the loading factors are chosen so as to maximise her expected profits. Figure 5 reveals that the insurer earns higher expected profits from insuring the two risks separately than from bundling across all simulated values of γ. Moreover, the size EP_{diff} of the insurer's preference toward separating over bundling is monotone increasing in γ, which implies that supplying separate contracts becomes increasingly profitable relative to bundled contracts as the consumer becomes more risk averse. That both the consumer and insurer earn excess surplus from separating when $\lambda_b = \lambda_b^*$ and $\lambda_s = \lambda_s^*$ implies that separating is socially optimal when the insurer is a monopoly. This yields another explanation for the abundance of separate contracts in

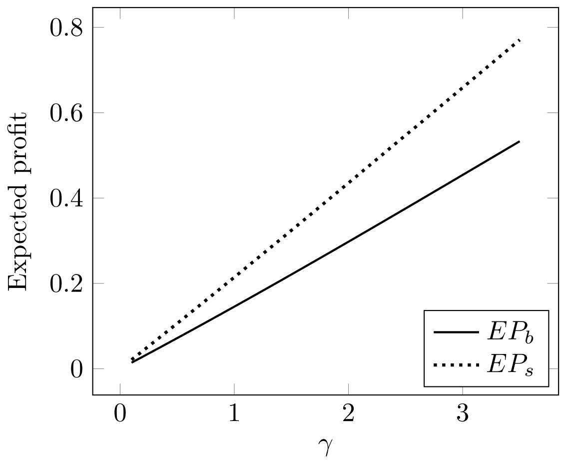

Figure 5: Insurer's expected profit with $\lambda_b = \lambda_b^*$ and $\lambda_s = \lambda_s^*$

real-world insurance markets: insurers may be exercising (approximate) monopoly power.

4.3 Asymmetric Information with Profit-Maximising Loading Factors

In section 4.1, we saw that if the insurer can transfer wealth to the consumer in exchange for his purchasing separate insurance contracts rather than the preferred bundling contract then there is, for sufficiently large values of γ, a Kaldor-Hicks transfer such that, upon the receipt of said transfer, the consumer will prefer to insure the two risks separately. In section 4.2, we saw that if the insurer has perfect information then no such transfer is needed because the relative attractiveness of buying separate contracts is endogenously included via the profit-maximising loading factors. Our final study relaxes the assumption of perfect information. In contrast to the standard theory of asymmetric information, in which the insurer is assumed to not know the probabilities of the risks of each consumer, we concentrate our analysis on the consumer's risk aversion. We retain our assumption that the stochastic structure of the two risks being insured are perfectly known by the consumer and insurer.

Our study proceeds as follows. The insurer observes that the market is composed of two types of consumers: those with high risk aversion, indicated by their having $\gamma = 3$, and those with low risk aversion, indicated by their having $\gamma = 1$. She also knows that the proportion of high risk aversion-types within the market is represented by some known parameter α with $0 \leq \alpha \leq 1$. For a given value of α, the insurer considers two options:

1. Offer only a pooling contract in which a single loading factor on separate insurance is offered to the entire market.

2. Offer a non-pooling menu of contracts consisting of a bundled contract and pair of separate contracts, each with a single loading factor.

As in the previous section, the insurer chooses a loading factor on each contract offering and consumers respond by choosing their expected utility-maximising deductibles. Under the non-pooled option, the separate contracts are intended for high risk aversion-types while the bundled contract is intended for low risk aversion-types. This choice of contract targeting is due to the fact that, as seen in Figure 5, the insurer earns greater expected profits from supplying separate contracts and from selling to consumers with high risk aversion. Offering a bundled contract allows the insurer to extract profits from low risk aversion-types. She cannot offer another separate pair of contracts for those types because the associated loading factor would need to be lower to offset their lower willingness to pay. However, lowering this loading factor would attract the high risk aversion-types away from their intended contract. We therefore place two incentive compatability constraints upon the choice of non-pooled loading factors, respectively ensuring that the high risk aversion-types and low risk aversion-types prefer the contract that is designed for them rather than that designed for the other type.

We compute the insurer's expected profit from the pooling option as

$$EP_p = \alpha EP_s^{\text{high}} + (1-\alpha)EP_s^{\text{low}},$$

where EP_s^{high} is her expected profit from selling a profit-maximising separate contract to a consumer with $\gamma = 3$ and EP_s^{low} is her expected profit from selling a profit-maximising separate contract to a consumer with $\gamma = 1$. Similarly, the insurer's expected profit from the non-pooling option is given by

$$EP_n = \alpha EP_s^{\text{high}} + (1-\alpha)EP_b^{\text{low}},$$

where EP_s^{high} is her expected profit from selling a profit-maximising and incentive compatible pair of separate contracts to a consumer with $\gamma = 3$, and EP_b^{low} is her expected profit from selling a profit-maximising and incentive compatible bundled contract to a consumer with $\gamma = 1$. Figure 6 plots EP_p and EP_n for $0 \leq \alpha \leq 1$. The interesting feature of this plot is the existence of multiple switching points for what is the optimal strategy for the insurer. She prefers to pool when $\alpha \leq 0.3$. For $0.3 < \alpha \leq 0.65$, the insurer instead prefers the non-pooled option. From $\alpha = 0.65$ onwards, the insurer returns to prefering the pooled option before becoming indifferent at around $\alpha = 0.85$.

That the pooled option is preferred when $\alpha = 0$ follows directly from Figure 5, which shows that, for both risk aversion types, the profit-maximising option is to insure the consumers' two risks via separate contracts. When $\alpha = 0$, the market is composed only of low risk aversion-types and so the insurer should offer a pair of separate contracts with the profit-maximising loading factor λ_s^* computed in section 4.2 for $\gamma = 1$. Since this generates higher expected profit than selling a bundled contract to a consumer with $\gamma = 1$,

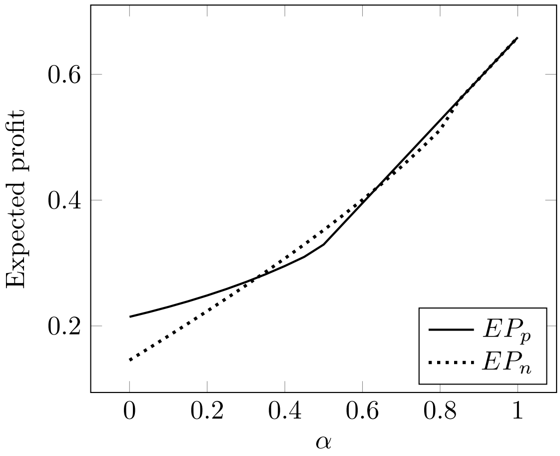

Figure 6: Insurer's expected profit under asymmetric information

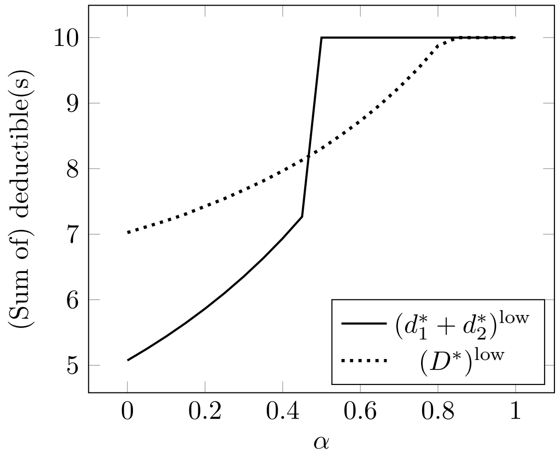

Figure 7: (Sum of) optimal deductible(s) when $\gamma = 1$

the pooled equilibrium is more profitable than the non-pooled equilibrium at $\alpha = 0$. As α grows, the profit from the non-pooled option grows and gradually takes over as being the better option for the insurer. At $\alpha = 0.5$, the pooled option becomes too expensive for low risk aversion-types. This can be seen in Figure 7, which plots the sum $(d_1^* + d_2^*)^{\text{low}}$ of the optimal separate deductibles chosen by the consumers with $\gamma = 1$ under the pooled option and the optimal bundled deductible $(D^*)^{\text{low}}$ chosen by those consumers under the non-pooled option. That $(d_1^* + d_2^*)^{\text{low}} = 10$ when $\alpha \geq 0.5$ means that the low risk aversion-types buy no insurance under the pooled option because their expected utility-maximising pair of separate contracts pay a zero indemnity in all states. Hence, for all values of α above 0.5, low risk aversion-types are only insured if the non-pooled option is available, in which case they buy a bundled contract. As a result, when the insurer's optimal strategy returns to being pooled the only consumers left in the market are the high risk aversion-types. Thus, for values of α beyond 0.65, the low risk aversion-types are absent from the market and the only insurance sold are the profit-maximising separate contracts to the high risk aversion-types.

The upshot from this analysis is that the dominant insurance product is separation rather than bundling. Separate insurance is the only product that is observed whenever the pooled option is favoured (for both low and high values of α). For the set of values of α over which the non-pooled option is preferred, we observe that for about half of that interval the value of α is larger than 0.5 and so there are more high risk aversion-types buying separate contracts than there are low risk aversion-types buying bundled contracts. Thus the observed tendency among real-world insurance markets to trade separate contracts can be justified theoretically by appealing to the idea that most consumers who demand insurance have a high level of risk aversion.

5 Conclusion

In this paper, we considered the question of whether independent risks should be bundled prior to seeking insurance or whether they should be insured separately. We have shown via numerical simulation that there certainly exist cases—and, we expect, a wide range of cases—in which the only product types that would be observed in the insurance market involve separate insurance of multiple independent risks rather than bundling those risks into a single contract. The clearest case in which this happens is when the insurer is able to set the loading factor on insurance products so as to maximise her expected profit and when she has perfect information with regard to the consumer's choice environment. In that case, and for both the power and exponential utility classes that we have considered, the unanimously preferred insurance product is separate contracts on the two risks.

We also considered a case in which there is asymmetric information on consumers' level of risk aversion within each utility class. In that case, the preferred insurance product depends critically on the fraction of insurance consumers who are high risk aversion-types. However, our simulations point towards the dominant insurance product being separate insurance for the two risks rather than bundling.

The present work is, for now, confined to a simulation exercise. Thus, while clearly illustrative of general principles, we cannot yet state any theorems that are unconditional upon the actual parameters used for the particular simulations that we have carried out. The one result that will carry over to any risk averse consumer is that fact that, due to Arrow's theorem, he will always prefer the bundled option to the separate option if both are priced with the same loading factor. However, we conjecture that at least one of our results will carry over more generally; namely the global preference for insuring the two risks separately when the insurer is able to set the contract loading factors at their profit-maximising values. This issue is left to our research agenda.

We also plan to develop our analysis of the case with asymmetric information, which has delivered some eye-opening results. Our findings from this analysis can be summarised as follows. For small values of α, the insurer's optimal strategy is to pursue a pooling equilibrium via separate contracts. For somewhat higher values of α, this optimal strategy switches to offering the non-pooled option with high risk aversion-types on the separate contracts and low risk aversion-types on the bundled contract. Finally, for even larger values of α, the only insurance that is sold is separated insurance to high risk aversion-types while the low risk aversion-types are left uninsured. We wonder if the switch from the pooled to non-pooled option is a general feature, or if there are cases in which the only insurance product ever bought is the pooled option with all consumers insuring their endowed risks via separate contracts.

As it stands, the only "comparative statics" type of analysis that is captured in our simulations is that of the consumer's level of risk aversion. However, we are also able to investigate the comparative statics features of the risks themselves. In the simulations that we performed, the two independent risks are exactly equal with each having the same probability distribution. It would certainly be of interest to pay closer attention to how the consumer's and insurer's preferences vary when the two risks are asymmetric.

6 References

Arrow, K. (1963). Uncertainty and the Welfare Economics of Medical Care. *American Economic Review 53*(5), 941–973.

Bakos, Y. and E. Brynjolfsson (1999). Bundling Information Goods: Pricing, Profits and Efficiency. *Management Science 45*(12), 1613–1630.

Evans, D. S. and M. A. Salinger (2008). The Role of Cost in Determining When Firms Offer Bundles. *Journal of Industrial Economics 56*(1), 143–168.

Koehl, P.-F. and B. Villeneuve (2001). Complementarity and Substitutability in Multiple-Risk Insurance Markets. *International Economic Review 42*(1), 245–266.

Sheikhzadeh, M. and E. Elahi (2013). Product Bundling: Impacts of Product Heterogeneity and Risk Considerations. *International Journal of Production Economics 144*(1), 209–222.

Thatcher, M. E. and E. K. Clemons (2000). Managing the Costs of Informational Privacy: Pure Bundling as a Strategy in the Individual Health Insurance Market. *Journal of Management Information Systems 17*(2), 29–57.